

Corpus analysis without prior linguistic knowledge - unsupervised mining of phrases and subphrase structure

Stefan Gerdjikov ^{*1,2} and Klaus U. Schulz ^{†1}

¹*CIS, Ludwig Maximilians University, Munich, Germany*

²*FMI, University of Sofia "St. Kliment Ohridski", Sofia, Bulgaria*

February 19, 2016

Abstract

When looking at the structure of natural language, “phrases” and “words” are central notions. We consider the problem of identifying such “meaningful subparts” of language of any length and underlying composition principles in a completely corpus-based and language-independent way without using any kind of prior linguistic knowledge. Unsupervised methods for identifying “phrases”, mining subphrase structure and finding words in a fully automated way are described. This can be considered as a step towards automatically computing a “general dictionary and grammar of the corpus”. We hope that in the long run variants of our approach turn out to be useful for other kind of sequence data as well, such as, e.g., speech, genom sequences, or music annotation. Even if we are not primarily interested in immediate applications, results obtained for a variety of languages show that our methods are interesting for many practical tasks in text mining, terminology extraction and lexicography, search engine technology, and related fields.

1 Introduction

While there is no complete agreement on what the structural subparts of natural language exactly are, there is a general understanding that natural language is built

^{*}st.gerdjikov@abv.bg

[†]schulz@cis.uni-muenchen.de

in a compositional way using characteristic classes of substrings (or utterances) such as phrases, words, and morphemes. In natural language processing (NLP) and related fields, all these units play an important practical role. Phrases are used, e.g., in Statistical Machine Translation (SMT) [ON04, Koe10], Text Classification [Seb01, DIK11], and Information Retrieval [MRS08, Lui07, ZLC⁺13]. Words (“tokens” and “types”) represent the basis for text indexing approaches in Information Retrieval, and since centuries linguists collect words in general and specific lexica, including typical contexts and collocations. Morphemes carry information on part-of-speech, syntax and semantic roles [Lui07].

In this paper we present results from an ongoing research project which is centered around the problem of *how to detect structural subparts of natural (or other) language* and is characterized by the following *methodological principles*.

1. Try to detect *meaningful* units of *any length*, as opposed to, e.g., fixed length *n*-grams of words or letters.
2. Only use *corpora (or sequence data)* as an *empirical basis* to find these units.
3. When analyzing corpora, *do not use any kind of prior linguistic knowledge* and *no knowledge on the nature/functionality of a particular symbol* (such as blank, hyphen, brackets, uppercase vs. lowercase etc.) but rather simple mathematical assumptions.
4. Proceed in a completely *unsupervised* way.

While we acknowledge the success of NLP approaches based on *n*-gram statistics, it is clear that using units of fixed length does not naturally lead to meaningful subparts. In anatomic terms, our aim here is to find the “muscles, organs, and nerves” of a given language corpus. The other above principles are related in the sense that they stress purely empirical evidence. Corpora, from our point, are just raw sequences of symbols. For analyzing corpora we do not use prior linguistic knowledge since we want to see what “the corpus structure itself tells us about language”. This might help to compare existing linguistic knowledge and hypotheses with “purely empirical” corpus-driven evidence. Furthermore, principles and methods obtained from this puristic approach can be applied to *any* (not necessarily natural) language and might contribute to the analysis of other kind of sequence data such as biological sequences, sequence of gestures, animal language, music annotation, and others. We avoid any kind of supervision since we are not primarily interested in the path from prior human linguistic knowledge to the analysis of data, but rather on the question how structure can be detected and then “des-tilled” to knowledge. The focus of this paper is the mining of phrases, subphrase

structure, and words. From a practical side, one vision is to automatically and “online” create a kind of “general dictionary and grammar of a given corpus” that collects, connects and relates all interesting language units of the corpus, with a direct reference to all occurrences. In the Conclusion we add a remark on possible interaction of our approach with unsupervised morphology induction approaches, such as [Gol01, Zem08, CL05, SO15].

As a matter of fact, this programmatic approach is by no means new. As an early influential advocate, Zellig Harris [Har54, Har91] has stressed that corpus analysis based on mathematical principles should be used to derive empirical evidence for linguistic hypotheses. Later, Maurice Gross [Gro93, Gro99], Franz Guenther [Gue05], Johannes Goller [Gol10] and others followed this line. Corpus analysis is often used to find lexico-semantic relations [Sto10]. Further related work is described in Section 6. One algorithmic contribution of our work is the use of a special index structure for corpora as a basis: symmetric directed acyclic word graphs (SDAWGs) [IHS⁺01] represent a generalization of suffix trees [Ukk95]. Using the SDAWG index of a corpus, for any substring of the corpus we have immediate access to all *left and right contexts of arbitrary length*. In this way we see “how” a substring occurs in the corpus (i.e., in which contexts and larger sequences, how often), which provides an excellent basis for recognizing meaningful units and connected tasks.

When asking for the composition of natural language in terms of meaningful subparts, a fundamental question is how small parts are combined to form larger units. Most grammar approaches (e.g., context-free grammars) are “concatenation based” in the sense that always *disjoint, non-overlapping* parts are glued together to build larger units. In contrast, our base assumption is that in general “phrases” can overlap, and combination principles go beyond plain concatenation.

Our base view of language composition is briefly explained in Section 2. Using this view, Section 3 introduces mathematical principles and an algorithm for recognizing meaningful units and decomposing sentences into “phrases”. We illustrate the sentence decompositions and “phrases” obtained for a variety of distinct languages (English, German, French, Spanish, Dutch, Chinese) and corpora (including Wikipedia, Europarl, Medline). Section 4 introduces principles for further decomposing “phrases”. The base units found in this way (in most cases) are “words”. Recall that we start from a point where we do not have the concept of a “word”. Again we presents results obtained for distinct corpora and languages. In Section 5 we discuss some possible applications of the phrase detection mechanisms: first we show how subphrase structure can be used to find a ranked spectrum of “most characteristic words” of the given corpus, and to extract paradigmatic word networks based on co-occurrence of terms in phrases. We also look at stylis-

tic corpus analysis, terminology extraction, and automated query expansion for search engines. In Section 6 we briefly comment on related work. Section 7 gives a conclusion and comments on future work.

2 Phrase model - intuitions and motivation

As a starting point we look at the notion of a “phrase”. When we want to recognize “phrases” in a corpus-driven way we need as a starting point some primitive base assumptions how “meaningful phrases” can be recognized and distinguished from arbitrary, “meaningless subsegments” of language. In our view, recognition of phrases should proceed in parallel with the recognition of those language units that connect and interlink phrases. Our main motivation comes from psycholinguistics [HW03, SWC06, Tra07]: *Young children recognise function words but do not use them in the beginning*. This suggests that there are some natural principles that govern the structure of natural language that are strongly related to the use of function words. Children are able to detect these regularities, and to distinguish between language parts that serve to *combine and connect* meaningful parts - function words - on the one-hand side, and *meaningful* parts - content words and phrases - on the other side.

Function words, or closed-class words, alone bear little sense, [Fri52]. Their main function is to bind other words or expressions in a grammatically correct way to form more complex meaningful language units. To suit these needs of the language, only a few function words, e.g. prepositions, articles, suffice. Thus, although some specific word may act as a function word in a special domain, generally the set of function words is stable in every particular language.

In our situation we do not assume that we know the set of function words, and the notion of a “word” is partially misleading: we do not assume that “phrase connectives” are always words, we even do not have a notion of “word” at the beginning. Yet, for convenience, we say “function word” when we mean “phrase connective”. We assume that the corpus consists of sentences¹ which make sense. The remaining assumptions we make about phrases and function words in the corpus are:

(A1) Phrases occur in different contexts.

(A2) Sentences are built of (overlapping) phrases.

(A3) Phrases overlap on function words.

¹Other small meaningful text units like paragraphs also conform with our model.

The first assumption means that the corpus provides structural evidence that helps to recognize phrases. The circumstance that the corpus consists of sentences that make sense is important for the second and third assumptions². Assumption (A2) follows the intuition that the sense of a sentence is determined by the sense of the individual phrases and the way they are composed, whereas the final assumption reflects the fact that function words bind smaller meaningful language units to build larger ones. In our case, the smaller meaningful language units are the phrases we are looking for, and the larger ones are the sentences in the corpus.

Since function words alone, to a large extent, are lacking sense, we can fairly approximate the sense of the sentence with the sense accumulated by the individual phrases. Thus, we do not need to bother of modelling what *sense* is, rather we will be concerned with the issue how regular w.r.t. the given corpus the bindings between the phrases are. We try to summarise this intuition in the following two principles:

- (P1) The function words occur regularly as the boundaries of phrases in the corpus.
- (P2) A particular sentence should exhibit these regularities as good as its structure w.r.t. the corpus allows.

3 Computing phrases and function words

Following the considerations of the previous section, we first look at the problem of how to compute the phrases and function words of a corpus. In the first subsection we develop an algorithm to solve this task. Afterwards we present the results obtained for distinct languages and corpora.

3.1 Algorithmic principles

We consider a corpus, \mathcal{C} , consisting of a finite number of sentences. The individual sentences are merely nonempty strings $S_n \in \Sigma^+$ over a finite alphabet Σ . In practice, the alphabet Σ is the set of all the symbols, i.e. letters, punctuation, white space, etc., that were used to create the corpus. According to Assumption (A1), the phrases are among those *strings* $S \in \Sigma^*$ that occur in different contexts within the corpus \mathcal{C} . Formally, for a phrase S the following two properties must hold:

²However, this is not a prerequisite for the mathematical approach described in the next section.

1. there are distinct $a, b \in \Sigma$ such that aS and bS occur in \mathcal{C} , or S is a prefix of some of the sentences S_n , and
2. there are distinct $c, d \in \Sigma$ such that Sc and Sd occur in \mathcal{C} , or S is an suffix of some of the sentences S_n .

Strings with the above two properties are called **general phrase candidates** of the corpus, the set of general phrase candidates is denoted $\mathcal{G}(\mathcal{C})$. For example, if the strings `unfortunate_` and `_fortunately` occur as part of some sentences in the corpus³, then they witness that `fortunate` can be followed by at least two distinct letters, `l` and `_`, and similarly it can be preceded by at least two distinct letters, `n` and `_`. Therefore `fortunate` is a general phrase candidate in this case. However, if it turned out that the string `ortuna` is always preceded by the letter `f` in the corpus, then the string `ortuna` will be *not* a member of $\mathcal{G}(\mathcal{C})$. The same would be true if `ortuna` was always followed by a `t` in the corpus.

According to (A2), a sentence S_n in the corpus can be decomposed into a sequence of nonempty strings $(P_1, P_2, \dots, P_{k_n})$ with $k_n > 1$ such that $P_i \in \mathcal{G}(\mathcal{C})$. For instance, if a sentence starts like

`We_are_not_making_the_connection...`

it may be the case that $P_1 = \text{We_are_not_}$, $P_2 = \text{_are_not_making_the_}$, $P_3 = \text{_the_connection}$. In this case P_1 and P_2 overlap on the string $F_1 = \text{_are_not_}$, whereas P_2 and P_3 overlap on the string $F_2 = \text{_the_}$. In general, P_i and P_{i+1} overlap on a common possibly empty string $F_i \in \Sigma^*$, i.e. $P_i = P'_i \circ F_i$ and $P_{i+1} = F_i \circ P''_{i+1}$, such that:

$$S_n = P_1 \circ F_1^{-1} \circ P_2 \circ F_2^{-1} \dots \circ F_{k_n-1}^{-1} \circ P_{k_n}.$$

We call the F_i 's **function strings**. Upon termination of our approach will declare some of these strings to be *function words*. Recall that if P_i, P_{i+1} are in the set of general phrase candidates $\mathcal{G}(\mathcal{C})$, then each suffix of P_i satisfies Property 2 and each prefix of P_{i+1} satisfies Property 1. In particular, since F_i is a suffix of P_i and F_i is a prefix of P_{i+1} , it follows that each function string F_i is always a general phrase candidate.

In general, there can be many decompositions, $(P_1, P_2, \dots, P_{k_n})$, for a given sentence S_n . According to Principle (P2) we should select the one(s) that exhibit(s) the most regular sequence of function strings, $(F_1, F_2, \dots, F_{k_n-1})$. Further, according to Principle (P1), the function strings should occur regularly as the boundaries of phrases. We will incorporate these two principles in terms of appropriate probability and likelihood measures.

³We use the symbol “`_`” to make a white space explicit.

To do so, let us assume for a moment that we have already determined a multiset of phrases $\mathcal{P} \subseteq \mathcal{G}(\mathcal{C})$. Based on (P1) we are going to assess the property of a function string F to be a function word in the corpus. Next, based on (P2), we are going to determine the best decomposition of each sentence in the corpus as a sequence of general phrase candidates. This second step will allow us to refine (redefine) the multiset of phrases we have started with. We can then iterate this process until the multiset of phrases stabilizes. Based on the obtained probabilities and the eventual sentence decompositions we define the ultimate set of function words and phrases.

Formally, to assess the property of F to be a function string w.r.t. the current multiset of phrases \mathcal{P} , we simply consider how often F occurs as a proper prefix, suffix or proper infix in \mathcal{P} :

$$\begin{aligned} \text{pref}(F|\mathcal{P}) &= |\{F \circ S \in \mathcal{P} \mid S \in \Sigma^+\}| \\ \text{suff}(F|\mathcal{P}) &= |\{S \circ F \in \mathcal{P} \mid S \in \Sigma^+\}| \end{aligned}$$

$$\text{inf}(F|\mathcal{P}) = |\{S_1 F S_2 \in \mathcal{P} \setminus \{F\} \mid (S_1, S_2) \in (\Sigma^*)^2\}|.$$

This statistics accounts for multiplicities. For example, if $F = \text{the_}$ and the phrase $P = \text{the_connection_}$ occurs 12 times in \mathcal{P} , then P will contribute with 12 units to $\text{pref}(\text{the_}|\mathcal{P})$. Now, we define the probability of a string $F \in \mathcal{G}(\mathcal{C})$ to be a function word w.r.t. to a given multiset $\mathcal{P} \subseteq \mathcal{G}(\mathcal{C})$ as the empirical probability that F is both a proper prefix and a proper suffix of a string in \mathcal{P} :

$$p_{fw}(F|\mathcal{P}) = p_{pref}(F|\mathcal{P}) \cdot p_{suf}(F|\mathcal{P}) \quad (1)$$

where

$$p_{pref}(F|\mathcal{P}) = \frac{\text{pref}(F|\mathcal{P})}{\text{inf}(F|\mathcal{P})}, \quad p_{suf}(F|\mathcal{P}) = \frac{\text{suff}(F|\mathcal{P})}{\text{inf}(F|\mathcal{P})}.$$

To model (P2) we optimize the likelihood w.r.t. $p_{fw}(\cdot|\mathcal{P})$. This means, that given a sentence $S_n \in \mathcal{C}$ we search for those sequences of strings $(P_1, P_2, \dots, P_{k_n})$ ($k_n > 1$) whose sequence of overlaps $\bar{F} = (F_1, F_2, \dots, F_{k_n-1})$ maximises the function:

$$\ell(\bar{F}|S_n, \mathcal{P}, \mathcal{C}) = \prod_{i=1}^{k_n-1} p_{fw}(F_i|\mathcal{P}). \quad (2)$$

Determining the maxima of Equation 2 supplies us with (an) optimal decomposition(s) $(P_1, P_2, \dots, P_{k_n})$ for each sentence S_n in the corpus \mathcal{C} . We call a substring $S_n[i..j]$ of S_n an *optimal phrase candidate for S_n* if it belongs to some optimal decomposition for S_n . We call the overlap of two consecutive strings in an optimal decomposition for S_n an *optimal overlap*. Now, we are in a position to reconsider

the multiset of phrases \mathcal{P} we have started with. The intuition is, that if we have made a reasonable even if not perfect guess for the set of phrases in the first step, then we should have obtained good scores for the typical function words. Thus, in the second step, optimising the likelihood across the sentences the function words will guide us to optimal strings which better describe the proper phrases in \mathcal{C} . Following these considerations we define the new multisets

$$\begin{aligned}\mathcal{F}(\mathcal{P}) &:= \{S_n[j_1..j_2] \mid S_n[j_1..j_2] \text{ optimal overlap in a sentence } S_n\} \\ \mathcal{P}' &:= \{S_n[j_1..j_2] \mid S_n[j_1..j_2] \text{ optimal phrase candidate for a sentence } S_n\}.\end{aligned}$$

Although the approach described above reflects principles (A1) - (A3) as well as (P1) and (P2), there is one subtle point that it overlooks. The optimisation of the likelihood implicitly struggles for shorter sequences of phrases. For instance, if

$$\begin{aligned}P_1 &= \text{_are_not_making_the_} \\ P_2 &= \text{_the_connections} \\ P &= \text{_are_not_making_the_connections}\end{aligned}$$

are three general phrase candidates, the optimisation of Equation 2 in situations like `We_are_not_making_the_connections...` will always prefer the longer phrase, P , rather than the two shorter overlapping phrases P_1 and P_2 . Thus, $P = \text{_are_not_making_the_connections}$ contributes to the count $\text{inf}(\text{_the_}|\mathcal{P})$, but it will not contribute to the counts $\text{pref}(\text{_the_}|\mathcal{P})$ and $\text{suff}(\text{_the_}|\mathcal{P})$. This is a problem, because if `_are_not_making_the_` and `_the_connections` are indeed phrases, then the presence of the phrase

$$\text{_are_not_making_the_connections}$$

witnesses that `_the_` acts as binding element in this case. However, this example will contribute to *decreasing* instead of increasing the probability $p_{fw}(\text{_the_}|\mathcal{P})$.

To remedy this inconsistency in the model, we take additional care for the cases described above and reflect them in the probability measure for a function word. Specifically, we define the set of stable prefix (suffix) phrases of a phrase: let $P \in \mathcal{P}$. A prefix (suffix) P' of P is a *stable prefix (suffix) phrase of P* w.r.t. the set of phrases \mathcal{P} iff two conditions are satisfied:

1. P' is a phrase in \mathcal{P} ,
2. every prefix (suffix) P'' of P of length $|P''| > |P'|$ which is a general phrase candidate is again a phrase in \mathcal{P} .

By $SPref(P|\mathcal{P})$ and $SSuff(P|\mathcal{P})$ we respectively denote the set of all stable prefix (suffix) phrases of phrase P . Now, we account for the inconsistency of the model described above in the following way. We consider phrases $P \in \mathcal{P}$ which can be represented as two overlapping phrases $P_1, P_2 \in \mathcal{P}$ but at least one of them should be stable. This means that either P_1 is a stable prefix phrase of P , or P_2 is a stable suffix phrase of P . The *overlap value* $ov(F|\mathcal{P})$ of a function string F w.r.t. the set of phrases \mathcal{P} is defined as the number of phrases, P , in the form $P_1 \circ F^{-1} \circ P_2 \in \mathcal{P}$ where $(P_1, P_2) \in \mathcal{P}^2$ such that $P_1 \in SPref(P)$ or $P_2 \in SSuff(P)$. As before we account for possible multiplicities P in the set \mathcal{P} . The overlap value $ov(F|\mathcal{P})$ is used to boost the empirical probabilities $p_{pref}(F|\mathcal{P})$ and $p_{suf}(F|\mathcal{P})$:

$$\begin{aligned}\tilde{p}_{pref}(F|\mathcal{P}) &= p_{pref}(F|\mathcal{P}) + \eta_{pref}(F|\mathcal{P}) \frac{ov(F|\mathcal{P})}{inf(F|\mathcal{P})} \\ \tilde{p}_{suf}(F|\mathcal{P}) &= p_{suf}(F|\mathcal{P}) + \eta_{suf}(F|\mathcal{P}) \frac{ov(F|\mathcal{P})}{inf(F|\mathcal{P})}\end{aligned}$$

where $\eta_i(F|\mathcal{P}) = p_i(F|\mathcal{P}) / (p_{pref}(F|\mathcal{P}) + p_{suf}(F|\mathcal{P}))$ for $i \in \{pref, suf\}$. Essentially, the above formulae assume that the counts $pref(F|\mathcal{P})$ and $suff(F|\mathcal{P})$ are representative and thus the ratio $pref(F|\mathcal{P})/suff(F|\mathcal{P})$ reliably represents the true ratio of prefix-suffix property. Yet, due to cases accounted for in $ov(F|\mathcal{P})$, they are underestimated with respect to the total number of occurrences of F as an infix. This is why we distribute the amount $ov(F|\mathcal{P})$ as to maximise the entropy w.r.t. the empirical $pref(F|\mathcal{P})/suff(F|\mathcal{P})$. In the final model instead of p_{pref} and p_{suf} in Equation 1 we use \tilde{p}_{pref} and \tilde{p}_{suf} , respectively.

Algorithm for detecting phrases and function words. Our complete algorithm for detecting phrases and function words proceeds in the following way.

1. *Initialisation.* We start with the multiset of phrases $\mathcal{P}_0 := \mathcal{G}(\mathcal{C})$. The multiplicity of each string is simply the total number of occurrences of the string in the corpus.
2. *Iteration.* Afterwards we follow the spirit of the Maximum A Posteriori Principle, [BTS87, GJ08], and compute $\tilde{p}_{pref}^{(i)} = \tilde{p}_{pref}(\cdot|\mathcal{P}_i)$ and $\tilde{p}_{suf}^{(i)} = \tilde{p}_{suf}(\cdot|\mathcal{P}_i)$ and define the next set of phrases as $\mathcal{P}_{i+1} = \mathcal{P}'_i$ (see above).
3. *Termination.* As in other similar approaches the final goal is to obtain convergence. In our case we want to arrive at a state where the parameters \mathcal{P}_i , \tilde{p}_{pref} and \tilde{p}_{suf} stabilize. Since the probability functions \tilde{p}_{pref} and \tilde{p}_{suf} are entirely determined by the current set of phrases \mathcal{P}_i it is enough to define a convergence criterium for the multisets \mathcal{P}_i .

In order to assess the similarity between the multisets \mathcal{P}_i and \mathcal{P}_{i+1} , we use a pure set theoretical measure. Let $m_i(P)$ denote the multiplicity of P in the set \mathcal{P}_i ,

i.e. the number of times the general phrase candidate P is assigned to \mathcal{P}_i . We can express the size of the symmetric difference $\mathcal{P}_i \Delta \mathcal{P}_{i+1}$ of the two multisets \mathcal{P}_i and \mathcal{P}_{i+1} as

$$\|\mathcal{P}_i \Delta \mathcal{P}_{i+1}\| = \sum_P |m_i(P) - m_{i+1}(P)|.$$

Similarly we compute the size of the union $\mathcal{P}_i \cup \mathcal{P}_{i+1}$ of the two multisets as

$$\|\mathcal{P}_i \cup \mathcal{P}_{i+1}\| = \sum_P \max(m_i(P), m_{i+1}(P)).$$

As a similarity measure between the multisets \mathcal{P}_i and \mathcal{P}_{i+1} we use the relative size

$$\rho_i = \rho(\mathcal{P}_i, \mathcal{P}_{i+1}) = \frac{\|\mathcal{P}_i \Delta \mathcal{P}_{i+1}\|}{\|\mathcal{P}_i \cup \mathcal{P}_{i+1}\|}.$$

We iterate Step 2 until:

$$[\textit{Halting Criterion}] \quad \rho_i \leq \rho_{i+1} + \theta, \text{ where } \theta = 10^{-6}. \quad (3)$$

We will comment on this halting criterion at the end of this section.

4. *Final model.* Assume that the algorithm terminates after $n+1$ iterations of Step 2 when for the first time $\rho_{n+1} \geq \rho_n - \theta$. We define:

$$\begin{aligned} \mathcal{P} &= \mathcal{P}_n \\ \tilde{p}_{pref} &= \tilde{p}_{pref}^{(n)} \\ \tilde{p}_{suf} &= \tilde{p}_{suf}^{(n)} \\ p_{fw}(F | \mathcal{P}) &= \tilde{p}_{pref}(F | \mathcal{P}) \cdot \tilde{p}_{suf}(F | \mathcal{P}). \end{aligned}$$

We define the set of **function words of the corpus**, $\mathcal{F}(\mathcal{C})$, as the set of strings $F \in \mathcal{F}_n$ that satisfy the following properties:

- F has multiplicity at least $|\mathcal{C}|/10^3$ in $\mathcal{F}(\mathcal{P}) = \mathcal{F}(\mathcal{P}_n)$,
- $\tilde{p}_{pref}(F | \mathcal{P}) + \tilde{p}_{suf}(F | \mathcal{P}) > 0.4$ and
- $\tilde{p}_{pref}(F | \mathcal{P}) / \tilde{p}_{suf}(F | \mathcal{P}) \in (1/4; 4)$.

The first condition expresses that F not only occurs often in the corpus, but also serves as an overlap in many different cases. The second and third conditions⁴ capture the intuition that F is both a typical prefix and suffix boundary of a phrase, and not only, say, a typical prefix.

⁴Note that the inequality $\tilde{p}_{pref}(F | \mathcal{P}) + \tilde{p}_{suf}(F | \mathcal{P}) \leq 1$ is always fulfilled.

Our experiments have shown that the particular choice of these parameters is not significant. Particularly, among the strings satisfying the above conditions almost all achieve $\tilde{p}_{pref}(F|\mathcal{P}) + \tilde{p}_{suf}(F|\mathcal{P}) > 0.5$ whereas the ratio between the prefix and suffix property varies in spans $(1/3; 3)$. Additionally, typically the strings meeting these constraints have total probability at rates:

$$\tilde{p}_{pref}(F|\mathcal{P})\tilde{p}_{suf}(F|\mathcal{P}) > 9/100$$

where the possible maximum is $25/100$ (attained for $\tilde{p}_{pref}(F|\mathcal{P}) = \tilde{p}_{suf}(F|\mathcal{P}) = 0.5$). At the same the strings which are pruned have probabilities of one magnitude less.

We define the set of **phrases of the corpus**, $\mathcal{P}(\mathcal{C})$ as those $P \in \mathcal{P}$ that (i) start with some element of $\mathcal{F}(\mathcal{C})$ or are the beginning of a sentence and at the same time (ii) end with some element of $\mathcal{F}(\mathcal{C})$ or are the end of a sentence.

Remarks on the Halting Criterion. Naively, one could expect that the similarity ρ_i (normalized size of symmetric difference) between two consecutive sets of phrases converges to 0. However, in practice this does not happen. The reason for this phenomenon are rare strings which in particular sentences can serve both as overlaps in an ambiguous sequence of phrases. Thus, when they achieve good scores as overlaps the number of phrases they participate increases and this makes their scores drop in the next phase. However, when they are not active the part of the phrases where they are suffixes and prefixes become optimal and they boost the probabilities of these rare strings implying a vicious circle.

This is confirmed in our experiments, see Figure 1, where we depict in logarithmic scale the behaviour of ρ_i for distinct datasets. Multiset similarities ρ_i start with values around 1 and drop very quickly to values below 0.01. Afterwards values ρ_i slowly decrease to stabilise at rates 0.2×10^{-3} and 0.5×10^{-2} , but they do not continue to converge to 0. Nevertheless, the multisets \mathcal{P}_i exhibit a rather stable behaviour. In fact, more than 90% of the value ρ_i is due to cases where the multiplicity of a phrase P is increased or decreased by just 1, $|m_i(P) - m_{i+1}(P)| = 1$. Further the weighted *signed* symmetric difference between \mathcal{P}_i and \mathcal{P}_{i+1} :

$$\delta_i = \frac{\sum_P m_i(P) - m_{i+1}(P)}{\|\mathcal{P}_i \cup \mathcal{P}_{i+1}\|}.$$

ranges at rates $|\delta_i| \in [10^{-6}; 10^{-4}]$ and below, which shows that typically a “small gain” for one phrase is compensated by a “small loss” of another phrase. This means, that on average, the changes between $p_{pref}^{(i+1)}$ and $p_{pref}^{(i)}$ (resp. $p_{suf}^{(i+1)}$ and $p_{suf}^{(i)}$) are small for function words with high probabilities. This shows that function words with high probabilities $p_{fw}^{(i)}$ do not change significantly their probabilities to

$p_{fw}^{(i+1)}$. Whenever such a function word is admissible in a particular decomposition, it will be selected for a split. “Persistent turbulences” are caused by words occurring only once, or words occurring in a very specific syntactical contexts.

Technical realization. Technically, the set of all general phrase candidates $\mathcal{G}(\mathcal{C})$ can be represented after a linear preprocessing in linear space as a compressed acyclic word graph, [BBH⁺87, IHS⁺01], which generalises suffix trees, [Ukk95]. In this structure the required statistics, *pref*, *suff*, *inf* and *ov* can be easily collected by a straightforward traversal. Meanwhile, the optimisation required in Equation 2 reduces to a standard dynamic programming scheme. The graph necessary for this computation can be constructed by the means of an Aho-Corasick-like algorithm, [AC75]. Further speed-up can be achieved by an A-star algorithm, [HNR68, HNR72], built on top of the shortest path problem algorithm in proper interval graphs, [ACL93].

3.2 Experimental basis - languages and corpora

We considered corpora of different language, domain and size, see Table 1. The corpora cover six languages: English, German, French, Spanish, Dutch, and Chinese. Domains covered vary from philosophy through scientific medical abstracts, politics and journalistic articles, to general mixed topics. The size of the corpora scaled from several thousand to several Million sentences.

Wittgenstein Corpus. The Wittgenstein Corpus consists of 3871 original remarks of the German philosopher Ludwig Wittgenstein. Although the general language is German, the corpus also contains a few remarks in English.

Medline Corpus. The Medline Corpus is a collection of about 15 Million medical abstracts in English. Each abstract consists of about 10 sentences. We considered two subsets of this corpus. They were obtained by randomly selecting 0.5% and 10%, respectively, of the abstracts and gathering all the sentences in the resulting abstracts. The 0.5%-Medline Corpus consists of 800 K sentences, whereas the 10%-Medline Corpus consists of about 15 M sentences.

EU-Parliament Corpus. The EU-Parliament Corpus is a collection of the political statements in EU-Parliament from 1997 through 2012. The statements are maintained in all the languages of the EU-members at the date of the statement. Thus the entire corpus amounts to ca. 2 Million sentences for the languages of the long-lasting EU-members, English, German, etc. and to several hundreds of thousands for the more recent EU-members.

From this corpus we selected excerpts from the 2000-2001 sessions of the EU-Parliament in English, German, French, Spanish and Dutch. The size of the se-




Figure 1: Near-convergence of the algorithm for computing phrases. Behaviour of the similarity measure ρ_i for the symmetric difference of phrase multisets for distinct rounds i for corpora of different size, language and domain in logarithmic scale: — Accountant ($\sim 2M$ Dutch sentences, newspaper articles); — excerpt of EU-parliament ($\sim 20K$ German sentences, politics); — excerpt of EU-parliament ($\sim 20K$ French sentences, politics); — excerpt of EU-parliament ($\sim 20K$ English sentences, politics); — excerpt of EU-parliament ($\sim 20K$ Spanish sentences, politics); — excerpt of EU-parliament ($\sim 20K$ Dutch sentences, politics); — excerpt of Medline ($\sim 15M$ English sentences, medicine); — excerpt of Medline ($\sim 800K$ English sentences, medicine); — Wittgenstein ($\sim 4K$ German and English remarks, philosophy); (- - -) excerpt of the German Wikipedia Dump May 2014 ($\sim 2M$ sentences); (- - -) excerpt of the English Wikipedia Dump May 2014 ($\sim 2M$ sentences).

lected corpora amounts to 20 K sentences. We also processed the entire French and Spanish corpora, amounting at about 2 Million sentences, each.

Accountant Corpus. Accountant Corpus is a collection of the issues of the Dutch newspaper “Accountant”. The texts are from the period of World War I. The digital version of the corpus was obtained through OCR. We considered as *sentences* sequences of characters separated by at least two new lines. Single new lines were replaced by white spaces. In this way we tried to roughly reflect the natural reading order of paper articles. The obtained corpus, Accountant, contains about 500 K different sentences (of total 780 K) most of which have the structure of a paragraph or article.

Wikipedia. We considered also dumps of the English and German Wikipedias from May 14, 2014. In the English version we selected the pages AA – AG in their raw form. In the German Corpus we selected the pages AA – AH. We used WIKIPEDIAEXTRACTOR [AF13] to remove the hyperlinks. In both cases the sentences were defined as sequences of characters terminated by a new line. In this way we roughly reflect the paragraph structure of these electronic resources. The size of the English corpus amounted to 1,1 Million paragraphs of which more than 890 K different and the German Corpus resulted in about 650 K paragraphs, more than 510 K of which were distinct.

Chinese Corpus. “保镖天下” (“Bodyguard Legend”) is a Chinese electronic fiction which is freely available online for research, [htt]. It contains about 24K of small paragraphs or single sentences. The total number of characters is about 2M.

3.3 Results on function words

In this section we comment on the function words, $\mathcal{F}(\mathcal{C})$, that the algorithm described in the previous sections detected in the various corpora. Table 2 gives an overview for the European languages and distinct corpora. For each corpus we present the number of different strings in $|\mathcal{F}(\mathcal{C})|$, and we present a top segment of the function words. At this point, the order of the function words is determined by the number of times a particular string has occurred as an overlap in an optimal decomposition. The corresponding results for the Chinese corpus are illustrated at the end of this subsection.

Speaking in general terms, the number of the detected function words varies between 100–300. The results in Table 2 suggest that the actual number of function words does not depend on the size of the corpus, rather on its domain. For instance, the two Medline excerpts have a similar number of function words, although the size of one corpus is 20 times the size of the other. It is expected that in corpora with paragraph-like structure the number of function words increases. This is due

Corpus	Lang.	Domain	Size (sent.)	Size (words)
Wiki EN (May 2014)	EN	General	$\sim 1,1M$ (paragraphs)	$\sim 60M$
Wiki DE (May 2014)	DE	General	$\sim 650K$ (paragraphs)	$\sim 24M$
Accountant	NL	Newspaper articles	$\sim 780K$ (paragraphs)	$\sim 30M$
Medline 10%	EN	Medicine/Science	$\sim 15M$	$\sim 292M$
Medline 0.5%	EN	Medicine/Science	$\sim 800K$	$\sim 17M$
EU-Parl FR 100%	FR	Politics	$\sim 2M$	$\sim 52M$
EU-Parl ES 100%	FR	Politics	$\sim 2M$	$\sim 51M$
EU-Parl EN 1%	EN	Politics	$\sim 20K$	$\sim 535K$
EU-Parl DE 1%	DE	Politics	$\sim 20K$	$\sim 531K$
EU-Parl FR 1%	FR	Politics	$\sim 20K$	$\sim 630K$
EU-Parl ES 1%	ES	Politics	$\sim 20K$	$\sim 593K$
EU-Parl NL 1%	NL	Politics	$\sim 20K$	$\sim 577K$
Wittgenstein	DE+EN	Philosophy	$\sim 4K$ remarks	$\sim 250K$
Chinese	ZH	Fiction	$\sim 24K$ small paragraphs	$\sim 2M$ characters

Table 1: Metrics of the corpora considered in our experiments.

to binding words occurring in the beginning of sentences which in a sentence-based corpus could not be used as overlaps. The ratio of the corpus size to the number of corpus units also exerts certain role.

Further, in a domain-specific corpus, such as EU-Parliament or Medline, we may expect that certain words usually considered as content words will often occur at the boundaries of phrases. For example⁵, this is the case with `_patients_with_` in the larger Medline corpus and with `_patients_` and `_cell_` in the smaller one. A manual inspection of the results reveals that there are 10-20 such strings in the domain-specific corpora which have the above property, yet only one or two in the Wikipedia corpora, i.e. `_million_` and `_John_`.

Besides the particular cases mentioned above, the obtained lists comprise of white space, punctuation symbols, articles, prepositions, pronouns, auxiliary verbs. A significant number of function “words” are multi-tokens such as “to the”. All the function words in the Medline corpora and in most EU-Parliament corpora (EN, DE, NL and the smaller ES and FR) are delimited by white-spaces or punctuation symbols. There is a minimal number of exceptions in the other corpora (two in

⁵Here and in what follows we use the symbol `_` to highlight an occurrence of the white-space symbol.

the English Wikipedia, one in the larger Spanish, and three exceptions in the larger French corpus, five in the Dutch Accountant Corpus and German Wikipedia and two in the Wittgenstein Corpus). The reason for these exceptions, e.g. “e . _”, “es . _”, “s . _” in French or “en . _” in German (Wittgenstein Corpus) and Dutch (Accountant), “s ”_” in the English Wikipedia, can be explained in the following way. Some classes of words are underrepresented in the context of specific punctuation symbols. For instance, in German relative clauses often end with a verb with suffix *en*. Thus, we often observe “en . _” at the end of the sentence. Yet, considering a particular verb form it is by no means evident that it will occur twice at the end of a sentence.

Function words obtained for the “保镖天下” corpus. Since in Chinese the words of a sentence are not delimited it was interesting to see how the approach would work in this situation. Remarkably, results do not strongly differ from the other cases. For our Chinese corpus, “保镖天下”, which represents fiction, we obtained a total of 291 function words. We present the top 50 function words most often used in the decompositions with accompanying English equivalent. “” (empty string); “, ” (comma white space); “的” (of); “了” (over); “是” (is/are); “你” (you); “我” (I/me); “这” (this); “人” (people/human); “着” (no specific meaning in Chinese); “在” (at/in); “就” (no specific meaning in Chinese); “” (apostrophs); “黎箫” (Li Xiao – a person’s name in this fiction); “他” (he); “也” (also/too); “到” (arrive); “?” (?); “都” (both/all); “去” (go); “!” (!); “来” (come); “还” (still/again); “看着” (look/stare); “!” (!); “!” (!); “那” (that); “一个” (one); “出” (go/get out); “上” (up/above); “要” (need); “和” (and); “: ” (apostrophs) “自己” (myself); “将” (will/shall); “没有” (no/didn’t);

“能” (can/be able to); “后” (behind/later); “被” (be–used in passive tense, before a verb); “会” (will/can); “向” (towards); “已经” (already); “让” (let); “对” (write/correct); “却” (but/and yet/however); “给” (give/provide); “竟然” (unexpectedly/to one’s surprise); “我们” (we/us); “一脸” (the whole face); “现在” (now/currently);

3.4 Quantitative results on phrases

In this section we provide some quantitative statistics on the multiset of phrases \mathcal{P} extracted in the final step of our algorithm. From our point of view it is interesting to assess how complex the resulting phrases are and how often they have been selected by the algorithm to obtain an optimal decomposition of a sentence.

To assess the first characteristic of $\mathcal{P}(\mathcal{C})$ we counted in a *postprocessing* step the number of words, $word(P)$, contained in each of the strings $P \in \mathcal{P}(\mathcal{C})$. This

Table 2: Function words extracted from English, German, Dutch, French, and Spanish corpora: number of detected function words, $|\mathcal{F}(\mathcal{C})|$, and function words most often used in decompositions.

Corpus	$ \mathcal{F}(\mathcal{C}) $	function words used most often
WikiEN	300	-, the ; . ; and ; of ; in ; to ; a ; of the ; is ; " ; was ; . The ; in the ; for ; that ; by ; as ; with ; are ; - ; (; on ; were ; to the ; from ; at ; his ; . In ; " , ; ; an ; or ; which ;
EU-Parl (EN) 1%	198	-, the ; , ; to ; of ; a ; and ; in ; that ; be ; for ; this ; which ; in the ; are ; we ; will ; have ; on ; has ; an ; not ; to the ; with ; , the ; that the ; it ; as ; by ; to be ; , and ; more ;
Medline (10%)	145	-, of ; and ; the ; in ; of the ; to ; was ; with ; for ; a ; were ; is ; in the ; by ; that ; are ;) , ; from ; on ; at ; or ; after ; as ; +/- ; and the ; between ; during ; to the ; ; patients with ;
Medline (0.5%)	143	-, of ; the ; and ; in ; to ; of the ; was ; a ; were ; with ; for ; is ; in the ; by ; that ; are ; from ;) , ; at ; on ; as ; after ; to the ; or ; an ; patients ; +/- ; cells ; ; and the ;
WikiDE	302	-, . ; der ; die ; und ; des ; in ; von ; " ; - ; im ; das ; , ;) ; eine ; den ; als ; . Die ; zu ; mit ; ein ; ist ;) in der ; wurde ; für ; auf ; zur ; einer ; nach ; sind ; . Der ; nicht ;
EU-Parl 1% (DE)	158	-, , ; die ; der ; und ; zu ; in ; den ; eine , daß ; des , die ; das ; nicht ; von ; werden ; im ; ist ; ein ; für ; wir ; auf ; es ; einer ; auch ; zur ; mit ; dem ; sich ;
Wittgenstein ⁶	173	-, . ; [die ; der ; das ; [1] ; ? ; the ; & ; daß ; nicht ; eine ; ein ; ist ; den ; to ; es , ; zu ; ; a ; so ; von ; of ; that ; wenn ; was ; I ; er ; als ; einen ; wie ; aber ; sie ;
Accountant	373	-, de ; een ; en . . ; en ; in ; dat ; is , ; die ; . De ; (; niet ; op ; ; ; ; in de ; van het ; zijn ; voor ; ^ ; met ; als ; dat de ; . Het ; wordt ;
EU-Parl (NL) 1%	151	-, , ; van ; een ; het ; in ; en ; van de ; dat ; te ; op ; die ; is ; niet ; voor ; zijn ; de ; worden ; met ; om ; dat de ; deze ; heeft ; van het ; aan ; voor de ; ook ; er ; wij ;
EU-Parl (FR) 1%	145	-, , ; de ; des ; les ; à ; que ; et ; une ; un ; en ; du ; qui ; pour ; pas ; dans ; nous ; au ; a ; il ; aux ; plus ; que la ; sur ; se ; par ; être ; d ' ;
EU-Parl (FR) 100%	198	-, , ; de ; des ; et ; les ; la ; à ; que ; en ; pour ; une ; du ; qui ; dans ; un ; au ; est ; aux ; a ; á la ; sont ; que les , et , mais ; par ; ont ; sur ; d'une , les ; plus ; d'un ;
EU-Parl (ES) 1%	140	-, , ; de ; que ; en ; el ; y ; a ; los ; del ; una ; las ; un ; se ; para ; de los ; no ; con ; es ; al ; a la ; ha ; de las ; por ; la ; su ; que la ; más ;
EU-Parl (ES) 100%	172	-, , ; que ; la ; en ; y ; los ; el ; a ; para ; del ; una ; se ; un ; las ; con ; no ; al ; a la ; es ; por ; en la ; de las ; a los , que , y ; ha ; más ; su ; sobre ; como ; ; , pero ; y la ; a las ;

was done by counting the number of white spaces, $ws(P)$, in each of the strings $P \in \mathcal{P}(\mathcal{C})$ and subtracting one, unless the string itself turned out to be a beginning or an end of a sentences in the corpus. Thus, $ws(P)$ takes on values $0, 1, \dots$ and $word(P)$ takes on values $-1, 0, \dots$. The values $word(P) = -1$ and $word(P) = 0$ mean that the string contains no word. We found out that strings of this form are unavoidable, details are left out.

Figure 2 shows the distribution of the number of phrases in $\mathcal{P}(\mathcal{C})$ w.r.t. their $word(P)$ -lengths for $word(P) \leq 30$. The variety of longer phrases $word(P) = 2, 3, 4$ is larger than the variety of single-word phrases. It is worth to note that number of cases $word(P) = -1, 0$ is very small. Manual inspection of such cases shows that these strings result from some particular features of the corpus, e.g. special symbols such as “^” in the OCR-ed corpus “Accountant”, special nomenclature and indices of the remarks in the Wittgenstein Corpus, chemical formulas in the Medline Corpus etc.

In order to estimate phrases of what length were actually preferred by the algorithm we proceeded in the following way. For $w = -1, 0, \dots$ we computed the total number of multiplicities, $M(w)$, of a string $P \in \mathcal{P}(\mathcal{C})$ with $word(P) = w$, i.e.:

$$M(w) = \sum_{P \in \mathcal{C}: word(P)=w} m_n(P).$$

We also computed the total number of occurrences, $O(w)$, in the corpus of the strings $P \in \mathcal{P}$ with $word(P) = w$, i.e.:

$$O(w) = \sum_{P \in \mathcal{C}: word(P)=w, m_n(P) \geq 1} occ(P).$$

Thus, the ratio $M(w)/O(w)$ provides an average on which percentage of the occurrences of a phrase $P \in \mathcal{P}$ with $word(P) = w$ has been used in the final decompositions determined by \mathcal{P} . In Figure 3 we depict the relationship between w and $M(w)/O(w)$ for $w \leq 30$ for the corpora w.r.t. their size. Results confirm that the algorithm selects longer phrases with higher preference than shorter ones.

3.5 Qualitative results on phrases

In our approach, the set of phrases represents a subset of the set of general phrase candidates, $\mathcal{G}(\mathcal{C})$. The above conditions imposed on general phrase candidates imply that a string with only one occurrence in \mathcal{C} that does not represent a full sentence cannot belong to $\mathcal{G}(\mathcal{C})$. This means that only those proper phrases can be found that have at least two occurrences (with distinct immediate contexts) in the given corpus \mathcal{C} . Since long substrings often have just one occurrence, many long substrings




Figure 2: Distribution of the number of phrases containing n words: $word(P) \leq 30$ against the logarithmic scale of the number of phrases $P \in \mathcal{C}(\mathcal{P})$: Upper diagram: (—) excerpt of Medline Corpus ($\sim 15M$ sentences); (—) excerpt of Medline Corpus ($\sim 800K$ sentences); (—) Accountant Corpus ($\sim 2M$ sentences); (—) excerpt of German Wikipedia Dump May 2014 ($\sim 2M$ sentences, 5% of the entire Dump); (—) excerpt of English Wikipedia Dump May 2014 ($\sim 2M$ sentences, 2% of the entire Dump). Lower diagram: EU-Parliament Corpora excerpt ($\sim 20K$ sentences) in German (—); English (—); French (—); Dutch (—); Spanish (—); and the Wittgenstein Corpus ($\sim 4K$) remarks (—).




Figure 3: How often phrases are used for sentence decomposition - dependency on the number of words in the phrase: $w \leq 30$ (x-axis) and $M_C(w)/O_C(w)$ for the excerpt of Medline Corpus ($\sim 800K$ sentences) and the Accountant Corpus ($\sim 2M$ sentences) (upper diagram), and EU-Parliament Corpora excerpt ($\sim 20K$ sentences) in German (—); English (—); French (—); Dutch (—); Spanish (—); and the Wittgenstein Corpus ($\sim 4K$) remarks (—) (lower diagram).

that we would treat - from a linguistic point of view - as phrases are not found in $\mathcal{P}(\mathcal{C})$. As a matter of fact, we cannot obtain another result when we demand that phrases must occur in distinct contexts in the given corpus. From a qualitative point of view, the essential characteristics of the set of extracted phrases $\mathcal{P}(\mathcal{C})$ can be summarized in the following way. Examples for all cases are presented below.

1. Ignoring a small percentage of exceptions, elements $P \in \mathcal{P}(\mathcal{C})$ are sequences of “words” delimited by blanks or special sentence symbols (.,;:). At this point it should be kept in mind that our algorithm does not know the notion of a word, nor the role of distinct letters in the alphabet. Exceptions where elements $P \in \mathcal{P}(\mathcal{C})$ are not delimited by blanks (or sentence symbols) are caused by strings with only one occurrence in the corpus.
2. In most cases, elements $P \in \mathcal{P}(\mathcal{C})$ extracted are meaningful units in the sense that they represent phrases, or phrases extended on the left/right border by characteristic overlapping units, function words. In addition we often find units that would be interesting for language learners, showing how kernel sequences of words are typically connected with neighboured phrases. For the sake of reference, such units will be called *phrases extended by typical connectors*.
3. When ignoring delimiting functions words, linguistic phrases in the corpus in most cases are representable as elements in $\mathcal{P}(\mathcal{C})$ or sequences of such elements.

Example 3.1 We present the decomposition of the Europarl-sentence You have requested a debate on this subject in the course of the next few days , during this part-session . (en) and the parallel sentences for the languages French (fr), Spanish (es), Dutch (du), and German (de). In the following sentence decompositions, phrases $P \in \mathcal{P}(\mathcal{C})$ extracted are the substrings marked by upper brackets (...) or lower brackets (...), and overlaps of decompositions have the form (...), function “words” in overlaps highlighted in bold.

- (en) (You have(...)requested a(...)debate(**on this**)subject(**in the**)course of the next few days(), during this()part-session .)
- (fr) (Vous avez(...)souhaité()un débat à ce sujet()dans les prochains jours(),)au cours(**de cette**)période de session().)

- (es) (Sus Señorías han()solicitado un()debate sobre el
()tema para()los próximos días ,()en el curso(**de**)
este período()de sesiones .)
- (du) (U heeft aangegeven dat u()deze vergaderperiode()
een debat()wilt()over deze()rampen .)
- (de) (Im Parlament()besteht der()Wunsch nach(**einer**)
Aussprache ()im Verlauf()dieser Sitzungsperiode()
in den nächsten Tagen().)

Several long phrases with not just one occurrences in the corpus are recognized.
Examples are

in.the.course.of.the.next.few.days_
un.débat.à.ce.sujet_
in.den.nächsten.Tagen_

Looking at the sentence beginnings we find “phrases extended by typical connectors” in the above sense:

You.have_
Vous.avez_
U.heeft.aangegeven.dat.u_
Sus.Señorías.han_

Examples where long phrases can be represented as sequence of units in $\mathcal{P}(\mathcal{C})$
(possibly deleting overlaps on the two sides) are

You.have.requested.a.debate.on.this.subject
Vous.avez.souhaité.un.débat.à.ce.sujet
au.cours.de.cette.période.de.session_
de.este.período.de.sesiones_.
im.Verlauf.dieser.Sitzungsperiode_

Example 3.2 We present a typical decomposition of a sentence of the Medline Corpus.

(A(**viral**)etiology,()eg, the()congenital rubella syndrome,()was considered most likely,()but detailed() investigations(**proved to be**)negative.)

In this sentence, ignoring extensions caused by the comma “,”, terminological expressions such as

`_viral_etiology`
`_congenital_rubella_syndrome`

are found as phrases. In fact, since terminological expressions often appear in distinct contexts, one strength of the method is the ability to find terminological phrases (see Section 5.2). The noun phrase `detailed investigations` is split, at this point corpus characteristics cause the effect that `_but_detailed_` and `_investigations_proved_to_be_` are preferred. Phrases found also include

`_was_considered_most_likely,_`
`_proved_to_be_negative.`

Phrases that can be obtained by combining recognized units are, e.g.,

`_but_detailed_investigations_proved_to_be_negative.`

Also note that `_viral_` and `_proved_to_be_` occur as an overlap, a characteristic feature of the Medline corpus.

Example 3.3 We present a typical decomposition of a sentence of the English Wikipedia.

(She further()demonstrated that()much of the other() research in the book()arguing for()a link(**between**) Sir(**John**)Williams and()the Ripper crimes()was flawed.)

Here `_between_` and `_John_` occur as an overlap. Phrases found are

`_research_in_the_book_`
`_the_Ripper_crimes_`

Interesting units that language learners would consider as important expressions, but not representing linguistic phrases (“phrases extended by typical connectors”

in the above sense), are

_arguing_for_
_a_link_between_

Example 3.4 We present a typical decomposition of a sentence of the German Wikipedia.

(vom()Orchestre de la Suisse Romande)(, den()Berliner
Philharmoniker,()dem Tonhalle-Orchester Zürich()mit
Werken ()klass(ischer Komponisten), (wie)auch()Komponisten
() Polens(von)Frédéric Chopin()bis()Witold Lutoslawski
()über()Krzysztof Penderecki.)

Due to particularities of the corpus, `_ischer_Komponisten_` occurs as an overlap. Phrases directly found are

_Orchestre_de_la_Suisse_Romande
_Berliner_Philharmoniker,_
_dem_Tonhalle-Orchester_Zürich_

Phrases corresponding to combinations of recognized units are

_mit_Werken_klassischer_Komponisten_
_von_Frédéric_Chopin_bis_Witold_Lutoslawski_

Also note that person names are well recognized. This is again due to the effect that these person names appear in distinct contexts in the corpus.

Example 3.5 We present a typical decomposition of a sentence of the Chinese corpus. The sentence is

“老者顺手递过来一份文件，黎箫打开一看，眼睛立刻就直了，没看别的，就看那张不大的照片，美女啊，绝对的美女啊，我们黎箫口水“飞流直下三千尺””

the English translation is

“The old man handed over a document, LiXiao stared at it with motionless eyes as soon as he opened it, looking at nothing else, just the small picture, beautiful girl, a really beautiful girl, the saliva of our Lixiao has flowed three thousand feet.”

The sentence decomposition determined by our algorithm is the following.

“(老者)顺手(递过(来)一份(文件(,)黎箫打开(一看, 眼睛立刻就直了(,)没看(别的(,)就看(那张(不大(的)照片(,)美女)啊, 绝对的(美女啊(,)我们黎箫(口水(“飞(流(直下(三千(尺)”)

It has the following (partially overlapping) phrases:

“(老者old man)顺手smoothly(递过来hand (something to somebody)
(来一份one piece)(文件document, or paper,), Lixiao opens黎箫打
开(一看, 眼睛立刻looks, eyes immediately)就直了motionless/steady,)
(, 没看did not look(别的others,), 就看only look(那张that piece
of(不大的not big, small)(的照片picture,), 美女beautiful girl)啊,
绝对的absolute(美女啊beautiful girl,), 我们黎箫Lixiao or our Lixiao(口
水saliva) (“飞fly(流flow(直下straightly down(三千three thousand(尺”feet)”

3.6 Computing extended sets of phrases

An evident drawback of the approach proposed above is its incapability to detect phrases occurring only once: since these phrases do not appear in distinct contexts they are not contained in the initial set of general phrase candidates $\mathcal{G}(\mathcal{C})$. Obviously, many substrings occurring only once represent interesting language units. The problem is challenging because we do not have immediate statistical evidence that helps to detect such substrings. We looked at two ways to address this problem, in both cases using a dynamic approach.

1. Extended island phrases. We consider the sentences, S_n , in the natural order of their appearance in the corpus \mathcal{C} . At each time step, t , we let \mathcal{C}_t be the set of sentences observed up to moment t . Using a dynamic interpretation of assumptions (A1)-(A3), the set of phrase candidates $\mathcal{G}(\mathcal{C}_t)$ w.r.t. the restricted corpus \mathcal{C}_t only contains those strings that occur in different left and right contexts within \mathcal{C}_t . Now, when processing sentence S_{t+1} we compute “island” substrings V characterized by the following property:

1. V is a maximal infix of S_{t+1} which does not contain as a substring any phrase candidate in $\mathcal{G}(\mathcal{C}_t)$ bounded with function words, $\mathcal{F}(\mathcal{P})$.

“Island substrings” V with these properties are extended on both sides until we reach the sentence border or a delimiting function word. Extended island strings

\hat{V} are treated as candidates for phrases. For example, if sentence $S = S_{t+1}$ has the form

`The_Commission_proposal_is_a_start_but_it_is_not_enough.`

it might be the case “The_”, “_is_a_”, “_but_”, “_it_is_”, and “_not_” are already elements of $\mathcal{G}(\mathcal{C})_t$, but none of the remaining words in the sentence has this property. The strings “Commission_proposal”, “start” and “enough.” are islands not covered by these phrase candidates:

`The_Commission_proposal_is_a_start_but_it_is_not_enough.`

Assume that when decomposing the sentence these strings are decomposed in a brute force way into substrings like “C”, “ommission”, “pro”, “po”, “osal” and so on. None of these smaller strings is bounded by function words on both sides. In this situation we guess that

`_Commission_proposal_, _start_ and _enough.`

are phrases. Applying this simple procedure we extend the set of phrases, $\mathcal{P}(\mathcal{C})$ with these extended islands. Results obtained are illustrated in Table 3.

2. Instances of island schemes. In a second step we extended this idea. For each sentence S_t in the original corpus a copy S_t^a was introduced where the longest infixes which cannot be covered by phrases in \mathcal{P}_t are substituted by a new character “UNK”. For example, the above sentence would be transformed to

`The_UNK_is_a_UNK_but_it_is_not_UNK`

In this way we obtain a new corpus $\mathcal{C}^a = \{S_n^a\}_{n=1}^N$ consisting of abstracted sentence transformations of the original corpus. Abstract phrases detected in this corpus can be pulled back to \mathcal{C} , in this way recognising patterns that we were unable to detect originally. For example, if `The_UNK_is_a_UNK_` is a phrase in \mathcal{C}^a we may conclude that `The_Commission_proposal_is_a_start_` could be treated also as a phrase in the original corpus, \mathcal{C} . For phrase detection in \mathcal{C}^a we apply the algorithm described in Section 3.1. Afterwards we can simply take the resulting phrases containing UNK in \mathcal{P}^a - called *island schemes* - and replace the special symbol by the original island. Table 4 represents some island schemes and instances obtained by applying this procedure to the excerpts of English and German EU Parliament Corpus.

4 Computing subphrase structure and content words

In this section we further analyse the structure of phrases. As a side result we obtain a method for finding content words (non-function words) of the corpus.

implemented.without.delay
inaugural.speech
wilful.murder
medium-sized.businesses
Echternach.procession
asset.redistribution
displaying.intolerable.inertia
dogged.insistence
social.environment
workable.solution
bears.impressive.witness
non-consultant.hospital.doctors
road.accidents
modern-day.slaves
unfair.advantages
tatenlos.zusehen
bromierte.Flammschutzmitteln
mehrjährige.Ausrichtungsprogramm
gewisse.Entwicklungsrückstände.aufgeholt
gutgemeinte.Ratschläge
falschen.Illusionen.wecken
angestrebten.Programmplanungsziele
Reformmaßnahmen.eingeleitet
geleisteten.Wirtschaftshilfe
hemmungslosen.Subventionswettbewerb
funktionierenden.Wettbewerbsordnung
monatelange.Renovierung
rotgrüne.Bundesregierung
Wasserverschmutzung.ankämpfen
aussichtsreichsten.Anwärter
Einzelfall.herauszugreifen
Strukturfondsmittel.eingehen
zynische.Doppelzüngigkeit

Table 3: “Extended island phrases” extracted from an excerpt of the English and German EU Parliament Corpus.

Sample phrase in \mathcal{C}^a	Reconstructed phrase in \mathcal{C}
_of_UNK_and_the_UNK_of_UNK_	_of_liability_and_the_payment_of_compensation_
_of_UNK_and_the_UNK_of_UNK_	_of_detention_and_the _criminalisation_of_asylum-seekers_
_of_UNK_the_trust_of_	_of_eroding_the_trust_of_
_of_UNK_the_trust_of_	_of_regaining_the_trust_of_
_UNK_to_the_UNK_of_the_	_pay_particular_attention_to_the_wording_of_the_
_UNK_to_the_UNK_of_the_	_looks_forward_to_the_enactment_of_the_
_UNK_to_the_UNK_of_the_	_victim_to_the_rationalisations_of_the_
_UNK_in_terms_of_UNK_	_measured_in_terms_of_gross_domestic_product_
_UNK_in_terms_of_UNK_	_extensive_in_terms_of_compliance_
_UNK_in_terms_of_UNK_	_benefits_in_terms_of_welfare_
_UNK_of_UNK_which_	_phenomenon_of_solidarity_which_
_UNK_of_UNK_which_	_canons_of_beauty_which_
_UNK_of_UNK_which_	_rock_of_Sisyphus_which_
_UNK_zur_UNK_der_UNK	_integrierten_Strategien_zur _Wiederbelebung_der_Beziehungen_
_UNK_zur_UNK_der_UNK	_Zuverlässigkeitserklärung_zur_Ordnungsmäßigkeit _der_zugrundeliegenden_Vorgänge_
_in_den_UNK_und_UNK_	_in_den_Medien-_und_Kulturnetzen_
_in_den_UNK_und_UNK_	_in_den_Entbindungsheimen_und_Krankenhäusern_
_UNK_des_UNK_der_Union_	_Konsolidierung_des_Wirkens_der_Union_
_UNK_des_UNK_der_Union_	_Identitätsfaktor_des_Bildes_der_Union_
_UNK_des_europäischen_UNK_	_Büro_des_europäischen_Bürgerbeauftragten_
_UNK_des_europäischen_UNK_	_Teilung_des_europäischen_Kontinents_
_UNK_des_europäischen_UNK_	_Symbole_des_europäischen_Wissens_
_UNK_des_europäischen_UNK_	_Kernstück_des_europäischen_Landwirtschaftsmodells_
_in_UNK_Weise_UNK_	_in_lobenswerter_Weise_zurückgehalten_
_in_UNK_Weise_UNK_	_in_bemerkenswerter_Weise_gemeistert_
_in_UNK_Weise_UNK_	_in_unannehmbarer_Weise_legitimiert_

Table 4: Automatically extracted “island schemes” and instantiations in excerpts of the English and German EU Parliament Corpora.

4.1 Algorithmic principles

Now that we have determined a set of function words $\mathcal{F}(\mathcal{C})$ and a multiset of phrases $\mathcal{P}(\mathcal{C})$ we proceed to structure phrases in a hierarchical way. We address this issue in two steps. In the first step we explain how to generate the constituents of the phrases in $\mathcal{P}(\mathcal{C})$, which we call *subphrases*. As a by-product of this step we will be able to derive a notion corresponding to the common “words”. Recall that our algorithm just analyzes plain sequences of symbols without having any notion of word “built in”. In a second step, we continue to structure the resulting subphrases in a hierarchical way, obtaining a form of decomposition tree.

1. Splitting phrases and subphrases. Following the discussion in Section 2 and the model that we developed in Section 3, given a phrase $P \in \mathcal{P}(\mathcal{C})$ we look for a decomposition into substrings of $\mathcal{G}(\mathcal{C})$ and determine those that exhibit the most regular function words. Formally, we search for a decomposition of P into a sequence P_1, P_2, \dots, P_k of general phrase candidates P_i such that the sequence of induced function words $\bar{F} = (F_1, F_2, \dots, F_{k-1})$ optimises the likelihood:

$$\ell(\bar{F}|\mathcal{P}) = \prod_{i=1}^{k-1} \tilde{p}_{pref}(F_i|\mathcal{P}) \cdot \tilde{p}_{suf}(F_i|\mathcal{P}).$$

The only difference to the situation above is that we now have fixed the set of function words, requiring that $F_i \in \mathcal{F}(\mathcal{C})$ for all $1 \leq i < k$. Strings P_i only need to be general phrase candidates in $\mathcal{G}(\mathcal{C})$ but may fail to be elements of $\mathcal{P}(\mathcal{C})$. In this way, a larger spectrum of meaningful units in the corpus is obtained, beyond the multiset of (full) phrases in $\mathcal{P}(\mathcal{C})$. We call each P_i a **subphrase** in \mathcal{C} . Using the same principles again, each subphrase P can be further decomposed into finer subphrases until we arrive at **atomic subphrases** that cannot be further partitioned. By $\mathcal{SP}(\mathcal{C})$ we denote the union of $\mathcal{P}(\mathcal{C})$ with all subphrases obtained from iterated decomposition of phrases and subphrases. $\mathcal{SP}(\mathcal{C})$ is called the *set of subphrases of the corpus \mathcal{C}* . The set of atomic subphrases of \mathcal{C} is denoted $\mathcal{ASP}(\mathcal{C})$.

It is easy to see that each subphrase P that is not a sentence prefix (suffix) has a prefix (suffix) which is a function word in $\mathcal{F}(\mathcal{C})$. The **kernel** of P is obtained from P by deleting the prefix (suffix) of P in $\mathcal{F}(\mathcal{C})$ which has the maximal probability $p_{fw}(\cdot|\mathcal{P})$ as a function word. If P does not have such a prefix (suffix), we “delete” from P the empty prefix (suffix). In this way, from the set of phrases $\mathcal{P}(\mathcal{C})$ we obtain the set of **phrase kernels** $\mathcal{P}_K(\mathcal{C})$, from the set of subphrases we obtain the set of **subphrase kernels** $\mathcal{SP}_K(\mathcal{C})$, and from the set of atomic subphrases we obtain the set of **atomic subphrase kernels** $\mathcal{ASP}_K(\mathcal{C})$. It turns out that the atomic subphrase kernels of a phrase P typically are the “content words” contained in P , i.e., words that are not in the set of $\mathcal{F}(\mathcal{C})$.

2. Decomposition trees and functional schemes. For a phrase P we may select an optimal decomposition. When we continue to decompose the subphrases obtained, always selecting an optimal decomposition, we obtain a **decomposition tree** t for P . Each node of t represent a subphrase of P . The children of a node P' in t are obtained using the set of function words $\mathcal{F}(\mathcal{C})$ and the above optimization principle, selecting one optimal decomposition. A parallel **kernel decomposition tree** t_K is obtained replacing each node in t by its kernel. Using the two trees and looking at the leaves, each phrase P has a representation as a sequence of function words and atomic kernels. For phrases P that do not correspond to a sentence prefix or suffix, P is described as

$$P = F_0 A_1 F_1 \dots F_{n-1} A_n F_n$$

where $F_i \in \mathcal{F}(\mathcal{C})$ and the strings A_i represent the leaves of t_K . In this situation, the sequence $F_0 | F_1 | \dots | F_n$ is called a **functional scheme** for P . For sentence prefixes (suffixes) function word F_0 (resp. F_n) has to be omitted. A functional scheme can be considered as a rudimentary form of grammar rule. In general a (sub)phrase may have several optimal decompositions. Hence, for a given phrase P we obtain a **set of decomposition trees** $T(P)$, a parallel **set of kernel decomposition trees** $T_K(P)$, and a **set of functional schemes** $FS(P)$. However, we found that the functional scheme for a phrase and the leaf representation $P = F_0 A_1 F_1 \dots F_{n-1} A_n F_n$ is unique in most cases. From the multiset of phrases of the corpus \mathcal{C} we obtain the **multiset $FS(\mathcal{C})$ of all functional schemes** in \mathcal{C} , which may be ordered by the number of occurrences.

4.2 Results for distinct languages and corpora

As an illustration, in Figure 4 we present a decomposition tree and the parallel kernel decomposition tree for the German phrase

über den Vorschlag für eine Richtlinie des
Europäischen Parlaments und des Rates

Empty nodes in the kernel decomposition tree are caused by phrases that become empty when removing the function word borders. The functional scheme for the phrase has the form

| den | für eine | des || und des |.

Atomic subphrase kernels obtained are über, Vorschlag, Richtlinie, Europäischen, Parlaments, and Rates.

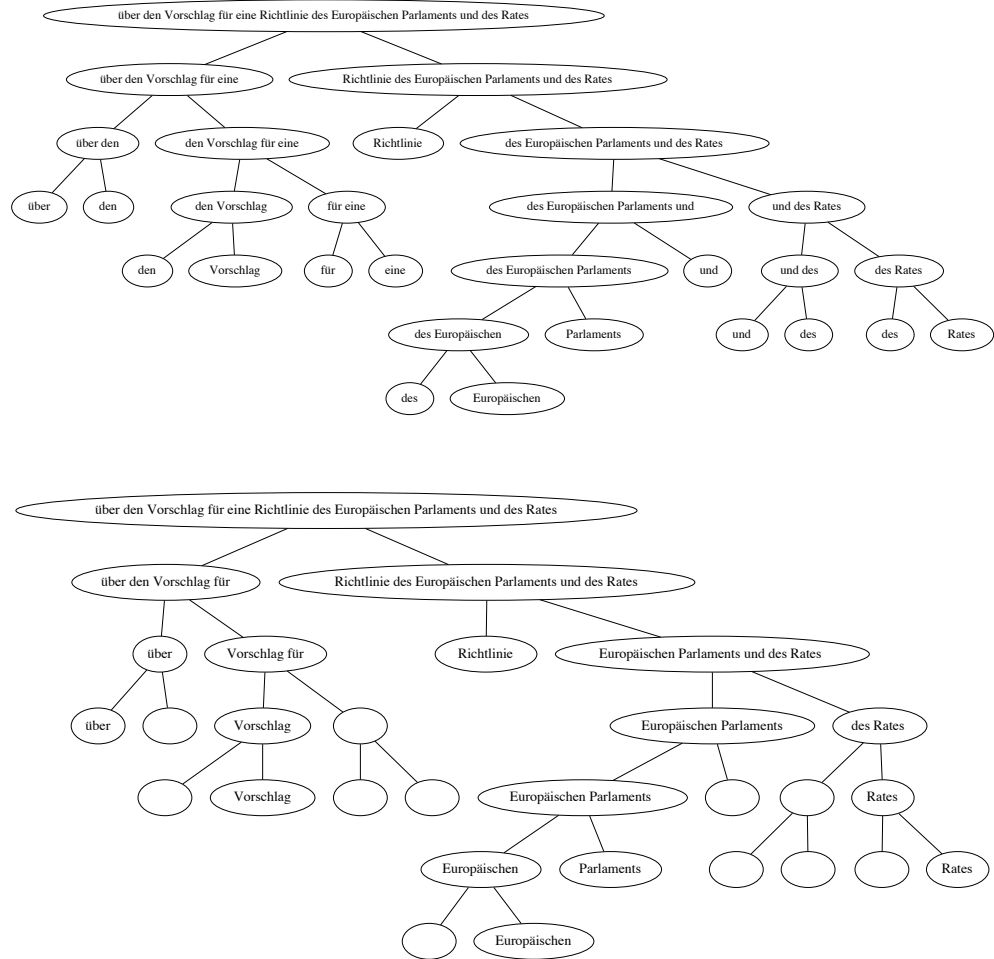


Figure 4: German decomposition tree and the parallel kernel decomposition tree.

The results for phrase decomposition in European languages are similar. As mentioned above, in almost each case, the functional scheme for a phrase turned out to be unique, even if there are several decomposition trees. This shows that the difference between distinct decompositions is mainly an issue of distinct orders in which to decompose (sub)phrases, not an issue of distinct final decomposition parts.

For all European languages and corpora considered, the set of atomic subphrase kernels obtained is essentially a lexicon of the content words of the corpus. Hence, despite of the fact that we do make any distinction between the distinct symbols in the textual alphabet, words are recognized as the “atomic” units from phrase decomposition. For Chinese, iterated subphrase splitting often leads to single symbols. Though multi-symbol words in general occur at one level of the decomposition of phrases, we yet do not have good principles to exactly find this point.

An obvious question is if the functional schemes obtained can be considered as a kind of induced grammar. However, there is no direct correspondence between traditional linguistic categories and the spectrum of functional schemes. Many of the functional schemes capture small sequence of words of a similar form with many instances, but the distinct instances do not always have a common linguistic functionality. In Table 5 we show the most frequent functional schemes for the English and German Wikipedia Corpora, respectively.

On the other hand we also found a considerable number of interesting patterns that point to a special form of grammatical construction. In Table 6 we add a selection of interesting complex schemes with several instances.

5 Applications

In this section we consider a selection of possible applications of the above methods. In each case the main point is to show that the automated computation of phrases and subphrases of a corpus and the bidirectional structure of the index structure used offer interesting possibilities to approach known problems on new paths.

5.1 Corpus lexicology

The predominant perspective in lexicology and corpus linguistics is to use a repository of corpora for collecting and verifying knowledge about a language or language variety. Under this perspective, language is the central object of study, and

Frame Scheme	# Occ.	Instance of a phrase with such a frame
— —	3355652	show soon
the —	1684261	the instinct
— of	1011444	laser of
. —	864564	. Invariant
— and	689264	courier and
— the	566956	exalting the
— to	551944	adherence to
— in	543862	dissolves in
of —	523440	of civilization.
a —	477598	a tolerance
and —	474076	and César
to —	408089	to emigrate
the — —	387926	the approximate area
in —	339115	in Brazil,
the — of	301400	the friends of
— — —	269382	future business partner
— of the	268339	observer of the
of the —	245283	of the society
— as	206565	livings as
— is	193940	reproduction is
— —	2723007	identische Formen
. —	954597	. Transaktionen
der —	688544	der Generaldirektor
die —	499024	die Besatzungen
— der	470917	Gegeninduktionsspannung der
— und	566956	Dardanellen und
, —	443004	, Ausnahmen
— die	308887	beginnen die
und —	306985	und Betriebe
— — —	257187	gute zeitliche Auflösung
— in	239913	Arbeitsgruppen in
von —	211493	Tragödie von
den —	201745	den Willen
in —	195174	in Bockenheim
— des	193070	Ordnung des
von —	183836	von Zufallszahlen
des —	168460	des Domes,
im —	166382	im Frühling
. — —	131594	. Daher spielten
das —	130116	das Mittelschiff

Table 5: Most frequent phrase frames in the English and German Wikipedia Corpus.

Frame Scheme	# Occ.	Instance of a phrase with such a frame
in an — to —	400	in an attempt to use
was — to — the	400	was devised to allow the
a — from the	399	a signal from the
— of — of the	398	consisting of members of the
at its —	398	at its zenith
the — and the —	397	the science and the arts
. The — of a	397	. The length of a
a — between the	396	a gap between the
— that the — of	199	believe that the presence of
the — of other	199	the decisions of other
in the — of — —	199	in the form of compress air
that — only	99	that time only
the — — of — and —	99	the atomic bombings of Hiroshima and Nagasaki
the — of — that	99	the branch of philosophy that
— of — games	10	family of card games
— a — system	10	patented a telegraph system
— a different —	10	take a different approach,
— — —" is — — if and only if	5	The field "F" is algebraically closed if and only if
mit — %	400	mit 13,7 %
, der — von	399	, der Nachfolger von
— bis zum —	396	reicht bis zum Jahr
ein — auf	396	ein Monopol auf
gibt es in —	396	gibt es in Nordamerika
über die — —	395	über die Teilung Deutschlands
— — als die	395	wesentlich älter als die
für — — und —	199	für physikalische Chemie und Elektrochemie
in das — der	198	in das Grundsatzprogramm der
als — für den	198	als Ursache für den
— der — am	99	beschloss der Bundesrat am
für die — der —	99	für die Durchsetzung der Menschenrechte
nach dem — des — —	99	nach dem Ende des NS-Regimes
— der — der Stadt —	10	An der Spitze der Stadt stand
— und — — Menschen	10	Neandertaler und anatomisch modernen Menschen
— als — Kilometer	10	mehr als sechs Kilometer
wurde die — zu einem — der	6	Wurde die Stadt zu einem Zentrum der

Table 6: Interesting functional schemes derived from the English and German Wikipedia Corpora.

corpora are just a means to obtain insights about language, e.g., by collecting the vocabulary found in the corpora. There is second perspective, which in general finds less attention: here a given corpus is considered as an independent object of study and the goal is to obtain insights about the “*language of the corpus*”. In this subsection we follow the second perspective. The methods described below, as those discussed above, can be applied in an unsupervised and fully automated way. They can be considered as a preliminary step to dynamically generate in an “online” manner views that characterize the vocabulary and language of a given corpus in distinct ways. The long-term objective is to enable “travels” in the language of the corpus, to compare the languages of distinct corpora etc. We start with an study where we try to find the most characteristic vocabulary of a given corpus. Afterwards we look at connections between words, semantic fields and phrase nets and at language style.

Characteristic words of a corpus. The problem of how to find “important” terms of a corpus or document collection has found considerable attention (see Section 6). Since both very rare and very frequent terms in general are inappropriate, there is no simple solution. Here we suggest a method that takes the composition of sentences and phrases into account. In order to find characteristic words of the corpus \mathcal{C} we first try to find the most “characteristic” or “specific” child P_i of a phrase $P \in \mathcal{P}(\mathcal{C})$ in the decomposition trees $t \in T(P)$. As we see below, at this step “rare” language units are preferred. Afterwards, looking at those kernels in $\mathcal{ASP}_K(\mathcal{C})$ that appear as the kernels of many “most characteristic subphrases”, we derive a ranked set of “characteristic words” or “key words” of the corpus \mathcal{C} . At this step, terms that are too rare obtain a low score. In order to formalise the intuitive notion of “characteristic”, we argue probabilistically. The intuition is that we want to decrease the uncertainty to obtain a longer phrase P starting from a shorter subphrase P_i . Probabilistically, this means that we need to maximise the conditional probability $p(P|P_i)$ where P_i ranges over all subphrases of P . Now, clearly:

$$p(P|P_i) = \frac{p(P, P_i)}{p(P_i)}$$

and if we look at the probability for these events as the event that the particular strings *occur*, then it is obvious that $p(P, P_i) = p(P)$ because P_i is a substring of P . Thus, if we naively consider the probability p as empirical probability, then the probability $p(P|P_i)$ is simply the ratio of the occurrences of P_i that imply an occurrence of P , i.e.:

$$p(P|P_i) = \frac{occ(P)}{occ(P_i)}.$$

Practically, the maximal value is simply obtained by selecting the child P_i of P with the smallest number of occurrences $occ(P_i)$ in the corpus. In our experiments, if two children have the same number of occurrences we simply selected the leftmost child.

Using the above measure for specificity we assign to each subphrase $P \in \mathcal{SP}(\mathcal{C})$ its most specific child P' . We write $P' = msc(P)$. Since a given subphrase P' in general acts as the most specific child of several superphrases P we obtain a forest

$$\mathcal{MSS}_{\mathcal{C}} = (\mathcal{SP}(\mathcal{C}), msc)$$

called the **forest of most specific subphrases**. The roots of $\mathcal{MSS}_{\mathcal{C}}$ are the “characteristic” elements in the set of atomic subphrases, $\mathcal{ASP}(\mathcal{C})$. The nodes of $\mathcal{MSS}_{\mathcal{C}}$ are the elements of $\mathcal{SP}(\mathcal{C})$, leaves representing maximal phrases. As before, as an alternative view we may look at the modified forest $\mathcal{MSS}_{\mathcal{C}}^K$ where we replace each subphrase by its kernel. The roots of this forest, which represent “characteristic” content words, are called **characteristic atomic kernels**.

As a measure for the importance of an atomic phrase P in $\mathcal{ASP}(\mathcal{C})$ we may use the number $\rho_{\mathcal{C}}(P)$ of distinct phrases in $\mathcal{P}(\mathcal{C})$ that are descendants of P in the forest $\mathcal{MSS}_{\mathcal{C}}$. In the experiments described below we found that selecting atomic phrases (or their kernels) with high $\rho_{\mathcal{C}}$ -score in fact yields a characteristic profile for a given corpus.

In our experiments we computed a ranked list of characteristic atomic kernels for our fragment of the Medline corpus, following the above method. The 60 top-ranked kernels are shown in Table 7. For the sake of comparison, similar ranking list were computed for the English Europarl corpus and for 2% of the English Wikipedia. In Table 7, pairs (n/m) present Europarl ranks (n) and Wikipedia ranks (m). Entries “-” mean that the entry did not occur in the ranked list for the corpus. The ranked lists for the corpora, despite of the large variety of word types in the lists, are radically distinct. When taking the top 60 kernels from Medline, only three (“only”, “one”, “time”) occur among the top 60 words in the ranking list for the Europarl corpus, and only seven (“used”, “only”, “found”, “one”, “developed”, “use”, “time”) occur among the top 60 words in the ranking list for the Wikipedia corpus. This shows that top segments of the ranking lists yield an interesting “profile” for a given corpus, which might also be useful for classification tasks and for stylistic studies (see below).

Connections between words, semantic fields and phrase nets. Automated approaches for exploring the paradigmatics of lexical units typically look at the co-occurrence of terms in documents, paragraphs, sentences or fixed size neighbourhoods [SP93, Sto10]. Another, more “syntactic” view is obtained when looking

1. observed (2310/602)	21. one (17/3)	41. tumor (-/10737)
2. those (121/30)	22. similar (1341/213)	42. children (1001/177)
3. used (220/1)	23. performed (5841/192)	43. identified (1864/608)
4. only (9/5)	24. control (362/150)	44. use (151/6)
5. treatment (2223/1106)	25. lower (3088/550)	45. different (231/194)
6. groups (361/142)	26. demonstrated (1167/1900)	46. examined (4705/5826)
7. present (127/114)	27. disease (3185/1674)	47. normal (3079/1211)
8. obtained (3743/1141)	28. age (1224/896)	48. time (41/2)
9. increase (534/358)	29. cells (-/1123)	49. blood (-/1376)
10. increased (991/386)	30. data (1947/277)	50. cases (561/316)
11. reported (4701/340)	31. reduced (1340/773)	51. studied (4507/846)
12. found (719/8)	32. decreased (5109/3881)	52. specific (284/462)
13. results (723/432)	33. risk (556/1182)	53. subjects (1068/2110)
14. higher (1253/505)	34. changes (728/270)	54. decrease (3735/2992)
15. patients (4104/1897)	35. group (198/66)	55. protein (7516/1941)
16. significantly (6063/958)	36. previously (3324/676)	56. 3 (720/779)
17. patient (-/3414)	37. developed (633/26)	57. compared (2692/1199)
18. cancer (-/2935)	38. high (532/226)	58. women (276/199)
19. study (1271/134)	39. effective (416/522)	59. investigated (3658/5008)
20. expression (1333/1666)	40. factors (1920/1089)	60. evaluated (8548/7615)

Table 7: Ranked list of characteristic atomic kernels from Medline, parallel ranks in the Europarl corpus (EP) and in the Wikipedia (W).

at the co-occurrence of terms in *phrases*. To this end, the forest of most specific subphrases MSS_C can be used. We compute the network $G = (\mathcal{ASP}(\mathcal{C}), E)$ where E is the set of all edges of the form $e = (A_1, P, A_2)$ where $P \in \mathcal{P}(\mathcal{C})$ is a descendant of the root nodes A_i in MSS_C ($i = 1, 2$). In the graph G , two atomic phrases A_1 and A_2 are strongly interlinked if they co-occur in many full phrases. Note that phrases P can have variable length and we do not ask if A_1 and A_2 are direct neighbours in P . Figures 5 and 6 show some of these graphs. Networks of this kind, when visualized in an appropriate way, offer interesting possibilities for mining and exploring corpora and related corpus vocabulary. Figure 5 starts from four characteristic words from Table 7, “blood”, “children”, “risk”, and “tumor”. Presented are atomic phrases that co-occur with these words in phrases, and the phrases where the co-occurrence appeared. Figure 6 is obtained in a similar way from the Wittgenstein corpus, starting from “Erklärung”, “Bedeutung” and “Gebrauch”.

Mining language style. It is a well-known phenomenon that the style even of *official* language differs when looking at distinct communities, fields or domains. When entering a new community, even native speakers often have difficulties to “adapt” language style and to meet the typical phrases and constructions of the

domain. For foreign language learners this task is even more complex. Many phrases automatically extracted from corpora using the above methods reflect the typical style of the corpus domain. For mining “stylistic” phrases we may start from subphrases such as “ we ”, “ my ”, and “it is ” and then look at the full phrases containing these strings. These phrases exhibit another specific “face” for each corpus. Below we present phrases including the strings “ my ”, and “it is ” derived (a) from the Europarl corpus, and (b) from Medline. It is important to note that the phrases extracted may contain the given string at any position, such as in “as a conclusion *it is* suggested that”, or in “express *my* regret at the”.

“my” in Europarl: my wholehearted support - my compliments to - extend my special thanks to - my congratulations to - my colleague , Commissioner Wallström - In my opinion , this is the only way - In my view , this - close to my heart - to thank my colleagues on the committee for - I expressed my views on - It is my belief that - On behalf of my - my honour to present to you - my group welcomes this - express my regret at the ...

“my” in Medline: to my knowledge this is the first - my own experience - my personal experience - studies in my laboratory - my professional life - of my research career - my assessment - my findings suggest that - answer to my question - my colleagues and i have developed

“it is” in Europarl: it is extremely important - it is therefore necessary to - it is a question of - it is , however , - it is in this spirit that - it is especially important - it is high time that - it is indeed the case that - it is my belief that - it is worth pointing out - it is unacceptable that - it is not enough to ...

“it is” in Medline: from these results it is concluded that - it is suggested that - it is shown that - it is argued that - it is postulated that - however it is not known whether - therefore it is suggested that - it is tempting to speculate that - indicating that it is - it is important to understand - furthermore it is shown that - as a conclusion it is suggested that ...

5.2 Terminology extraction

Automated terminology extraction helps to build up terminological dictionaries and thesauri, it has often been studied in the literature (Section 6). Most methods suggested use characteristic sequences of POS, word association measures such as log likelihood, mutual information etc., and corpus comparison to find terminological expressions in a given corpus. Here we suggest a three step approach only based on the techniques described above and using distinct corpora. In Step 1 we compute the characteristic atomic kernels of phrases for the given corpus, only relying on the input corpus. In Step 2, given the set of characteristic atomic kernels we try to single out “terminological” kernels (e.g., “cancer”) using the ranked lists of characteristic atomic kernels obtained from another corpora for comparison. In Step 3, given a terminological kernel (“cancer”) we try to find terminological (multi-word)

expressions that contain this kernel as substring (“female lung cancer”, “prostate cancer”).

Step 1 - characteristic atomic kernels. Our starting point for searching kernels is the decomposition of phrases described in Section 4. We compute the ranked list of characteristic atomic kernels. In general there are much more truly “characteristic” kernels than “terminological” kernels. Even if there is hardly a general agreement of what “terminological” exactly means, typically, mainly noun phrases are considered as terminological expressions. The list in Table 7 show that words of many distinct types can be “characteristic” for a given corpus.

Step 2 - terminological kernels. When trying to filter out “terminological” kernels from the list of “characteristic” kernels we could of course use POS or other type of linguistic information. However, for the sake of consistency of the approach we do not use external or human linguistic knowledge here. The numbers found in Table 7 suggest that “terminological” kernels typically do not appear in the top segment of the list obtained for the Wikipedia corpus. If we ignore kernels with a rank below 1,000 in one of the two corpora used for comparison, from the Medline top 60 segment we obtain the following list of kernels

treatment (2223/1106), obtained (3743/1141), patients (4104/1897), patient (-/3414), cancer (-/2935), expression (1333/1666),
demonstrated (1167/1900), disease (3185/1674), cells (-/1123), decreased (5109/3881), factors (1920/1089), tumor (-/10737),
examined (4705/5826), normal (3079/1211), blood (-/1376), subjects (1068/2110), decrease (3735/2992), protein (7516/1941),
compared (2692/1199), investigated (3658/5008), valuated (8548/7615)

We still find entries that are usually not considered as “terminological”, but clearly a much better focus is obtained. With larger bounds, the percentage of non-terminological kernels can be further reduced, a suitable compromise between precision and recall needs to be defined for each application.

Step 3 - kernel expansion. For each of the remaining top-ranked phrase kernels we consider the list of all phrases that contain this kernel as a substring. In this way, kernels are expanded to the left, to the right or in both directions. Since our phrases typically come with function word delimiters we clean each set by eliminating left and right “borders” that appear in our list of function words (see above). Cleaned substrings are ranked by frequency of occurrence. As a result we obtain a set of multi-word expressions. The characteristic features of the terminological expression found in this way are the following

1. the *number of words* in these expressions is *not fixed*,
2. the *kernels* used may *appear at any position* in such a multi-word expression.

As an illustration, in Appendix A we present upper parts of the ranked lists obtained for the kernels “disease”, “syndrome”, and “inflammatory”. In most cases, “disease” occupies the last position of the phrase. Expressions can be long - some examples from the full list are “risk factor for cerebrovascular disease”, or “intestinal failure-associated liver disease”. Examples where “disease” is not at the last position are “disease-related modification of exposure” and “disease is known to be caused”. For “syndrom”, an example is “metabolic syndrome components”. In contrast, “inflammatory” typically is found at the beginning of an expression, as expected. Examples not following this rule are “proinflammatory cytokine” and “nonsteroidal anti-inflammatory drugs”. We did not use any kind of additional cosmetics (which would easily be possible in practice). From our point, this third step is working extremely well.

5.3 Automated query expansion for search engines and free context size answer browsing

Querying a large corpus can be a difficult task since it is not simple to estimate in advance the number of answers to a search query. If the query is too specific, the number of answers might be small or empty, in other cases the answer set becomes extremely large and a manual inspection is not possible. In order to support users, search engines in the Internet often present a selection of possible right expansions to a unspecific string. For example, when typing “John” in the search window of the English Wikipedia, some expansions such as “John McCain”, “John Lennon”, and others (s.b.) are suggested.

The index-based technology presented above can be used to automatically produce interesting expansions. A given short query is just treated as a kernel in the sense of Section 5.2. Just using Step 3 “kernel expansion” of the above procedure, interesting *right and left* expansions for a query are found. In order to judge how well this approach works we indexed the complete English Wikipedia⁷ lists of possible expansions for a number of short queries. Possible expansions were ranked by the number of occurrences. Topmost Expansion suggestions for “John”, “Peter”, “Blue” and “Brown” are presented in Appendix B. We then compared this with Wikipedia’s expansion suggestions. Main insights obtained are:

1. Wikipedia only suggests a small number of expansions. The list of possible expansions produced by our method is much larger.

⁷The corpus, Wikipedia dumps May 2014, used for these experiments was split into sentences by PUNKT, [KS06], and afterwards processed by our approach from Section 3.1.

2. Wikipedia only produces right expansions for a query. Our lists also include left or two-sided expansion suggestions.
3. As to recall, most of Wikipedia’s suggestions are also found in our lists. In most exceptional cases, variants of Wikipedia’s suggestions are found.

The third point is illustrated in Table 8 where we show the full list of Wikipedia’s expansion suggestions (June 18, 2015) for “John”, “Peter”, “Blue”, “Brown”, “Italian”, and “French”. We show which suggestions are also found using our index-based “kernel expansion” method, and where we found variants of Wikipedia’s suggestions. The recall is excellent. In many cases, distinct variant phrases of Wikipedia’s suggestions were found in addition. Some of these are mentioned in Table 8.

Free context size answer browsing. Once the user has selected one of the expansion suggestions (“John McCain”), Wikipedia leads her to the Wikipedia page for the expansion (https://en.wikipedia.org/wiki/John_McCain). The symmetric index technology can be used for another type of interaction, which we call “free context size answer browsing”. We present to the user a *two-sided* list of concordances that includes all occurrences of the expansion in the indexed corpus. Since in the index left and right contexts of arbitrary length are directly accessible, the *size of contexts shown can be selected and varied by the user*. In Figure 7 we see a user interface with this functionality embedded. In this search interface, after typing any string (here: “agen”), the user may select the number of hits and the size of the concordances, which includes both left and right contexts for a hit. Using the lower button, the size of contexts presented can be changed in real-time. This kind of search interface is used at a project at CIS, LMU, where we collaborate with philosophers to find new and innovative ways of how to access corpora in Digital Humanities. See <http://www.cis.uni-muenchen.de/dighum/research-group-co/index.html>.

6 Related work

The discussion of related work is separated into two parts. We first consider contributions that are related when looking at the general characteristics of the approach presented above. Afterwards we take a more application oriented view.

John McCain	yes (356)
John Lennon	yes (602)
John, King of England	yes (9)
John, Paul, George, and Ringo	yes (2), also variants found
John Tolkein	"J. R. R. Tolkien" found for kernel "J. "
John A. Mcdonald	"John A. Macdonald" found (62)
John Wayne Gacy	yes (27)
John Roberts	yes (140)
Peter Ilyich Tchaikovsky	yes (3) "Pyotr Ilyich Tchaikovsky" found for kernel "Pyotr"
Peterborough	yes (1098)
Peter Parker, the Spectacular Spider-Man	yes (8)
Peter Sellers	yes (187)
Peter the Great	yes (311)
Peter Jennings	yes (47)
Peter IV of Portugal	yes (3)
Peterloo Massacre	yes (5)
Blue	yes (3251)
Blue's Clues	yes (55)
Blue whale	yes (19)
Bluetooth	yes (358)
Blues	yes (1931)
Blue ray disc	no
Blue-green algae	yes (2)
Blue iguana	"Blue iguana," found (3)
Blues rock	yes (4)
Blue1	yes (8)
Brown	yes (6564)
Brown v. Board of Education	yes (32)
Brown Mackie College	yes (7)
Brownhills	yes (49)
Brown Dog affair	yes (3)
Browning Hi-Power	yes (11)
Brown gold	no
Brown dwarf	"brown dwarf" (84) found for kernel "brown"
Brown-grey	"brown-grey" grey(37)found for kernel "brown"
Browning M2	yes (5)
Italian	yes (5409)
Italianate architecture	yes (14)
Italians	yes (1137)
Italian language	yes (166)
Italian War of 1521-26	"the Italian War of 1521-1526," found (4)
Italian films of the 2010s	"Italian films" found
Italian ballet	yes (3)
Italian War of 1542-46	"the Italian War of 1542-46" found (2)
Italian American	yes (100)
Italian Democratic Socialist Party	yes (3), "the It..." found (9)
French	yes (10538)
French language	yes (336)
French spacing (English)	"French Spacing" found (3)
French fries	yes (107)
French Revolution	yes (149)
French and Indian War	yes (49)
French colonial empire	yes (10)
French Mandate for Syria and the Lebanon	yes (2), "the Fr.." found (5)
French cuisine	yes (78)
French people	yes (105)

Table 8: Query expansion suggestions for “John”, “Peter”, “Blue”, “Brown”, “Italian”, and “French” from the English Wikipedia. Entries “yes” indicates that the same expansion was found automatically after indexing Wikipedia using above Step 3 “kernel expansion”. The number of occurrences of a phrase is given in brackets.

Welcome to Search Wittgenstein [Start Introduction](#)

Type Something:

Current Request: agen

Found 730 results. (in documents with width 45, showing 15 results).

Aber man kann sagen & beton, daß sie es enthält.
berhaupt nichts, was mit einer Mehrheit von Lagen (bzw. Größen) zu tun hat, oder aber etwas
ichtbar, und doch wissen wir, daß wir sie tragen (gleichsam indirekt).
h nach der Richtigkeit dieses Satzes nicht fragen (sowenig, wie nach der Richtigkeit der Form
Daher kann sie nicht sagen , daß Alles fließt.
Und nun könnte man sagen // Und nun höre ich die Mathematiker s
Als könnte man also fragen // darüber reden// , ob der König oder das
sagen // Und nun höre ich die Mathematiker s
it des Satzes aufgehoben wird . Man könnte sagen //; daß man s
icht existiert". // Es hat keinen Sinn zu sagen : "ich zerschneide töte etwas, was nicht
Ich werde mir sagen : "vielleicht ist er zu Hause aufgehalten
Wäre es nicht sozus agen Pedanterie, die Disjunktion noch immer nicht
Sozus agen Regeln die mit verblaßter Tinte geschrieben

Figure 7: Free context size answer browsing.

General characteristics of approach

Finding phrases. When looking for phrases, we do not use linguistic knowledge and only consider structural properties of the corpus. Our approach is completely language independent. In many application fields, as a common standard approach phrases for simplicity are replaced by n -grams of words, [ON04, Koe10]. Additional grammatical features, e.g. part of speech [Lui07], syntactic [ZLC⁺13], or semantic markers [EWVN11] may be applied to find n -grams that exhibit phrase characteristics. For larger values of n , statistical uncertainty is growing. Deep learning approaches [CWB⁺11, LST13] try to resolve this issue by deriving regularities in smaller dimension spaces. However, the embedding of data in a smaller dimension space is achieved by extracting features from a fixed length context or from the entire sentence, [CWB⁺11], with no attempt to reflect the general structure of the corpus. Recent approaches in this area try to account for this drawback by more advanced network topology, [TSM15]. Due to the fixed size of n -grams considered at one time, and since n -grams start at any point in the sentence, the subphrase structure of sentences is not analyzed.

Grammar development and induction. When trying to analyze sentence structure, grammar-based parsing algorithms can be used [Tom13]. The sentence and phrase decomposition strategies described in this paper can be considered as a preliminary step towards corpus based, automated, and language independent grammar induction. While our approach leads from corpus structure to grammatical rules, in parsing, fixed grammar rules impose a “predefined” structure on the corpus. For most languages, a full and correct syntactic analysis of arbitrary sentences in unrestricted texts is not possible. For searching phrases of a particular type, partial/local grammars [Gro97] can achieve high precision and recall. Until today, the development of (full, partial, or local) grammars with high coverage and precision is an ambitious and difficult task that needs a high amount of human work. Furthermore, most methods we are aware of are language dependent. Language independent and automated construction steps are found in probabilistic context-free grammars [Sak10] and learning regular expressions [DLT04]. The latter problem is mainly of theoretical interest.

Sequence analysis without linguistic knowledge. The problem of detecting meaningful units in a text corpus can be considered also as a kind of segmentation problem. In [LK09] and in [GGJ09] the authors consider the problem of segmenting sequences of letters without blanks and punctuation into words. [GGJ09] is focused on the segmentation of children’s utterances transcriptions into words. The authors present a Bayesian approach combined with the Maximum A Posteriori Principle in order to detect the most probable boundaries of words within

a sequence. In particular, they use Dirichlet distribution to model the number of different words contained in the corpus and a probability for an end of a word.

The approach presented in [LK09] is a generic one and the authors apply it also to Chinese word segmentation, Machine Translation alignment and others. The model defines only a probability for a sequence of letters to belong together. In this measure a key role is played by a fine-tuned parameter that controls the length of sequences. Other parameters of the model are trained via Expectation Maximisation Principle.

An extensive work has been also done on supervised word segmentation in Chinese. The methods range from statistical approaches based on CRFS [PFM04] to Recurrent Neural Networks [CQZH15] and others.

There are two main differences between the works cited above and the current paper. Firstly, the problem of segmentation requires that each letter is assigned to exactly one word. In contrast, we allow overlaps between these units which represents the possibility of sharing small pieces of information which however have high frequency and can provide us with reliable statistics. Secondly, it does not seem that previous approaches reflect the structure of the corpus, rather it is the statistics which aggregates the statistics obtained from the different entries of the corpus. In the presented approach we use this structure in a two-fold manner: (i) to constrain the number of available strings and thus define their possible boundaries in terms of other strings; and (ii) to find the best fit of these shorter strings within a particular corpus' entry. This considerations allow us to avoid the explicit modelling of phrase/word types and phrase/word length.

Index-based analysis of corpora. In [Gol10] (see also [Gue05]), suffix arrays are used as a text index structure for mining local grammars. POS-info stored in a synchronized index structure enables a rich set of queries including textual and grammatical conditions. In many fields, suffix trees/arrays are used for various tasks, applications ranging from sequence analysis in bioinformatics [Gus97] to text-reuse in digital humanities. The index structure used for our approach, symmetric directed acyclic word graphs (SDAWGs) [IHS⁺01], has yet not found much attention. The power of this index relies on the fact that it gives access to left and right contexts of any length. This feature is used in many of the above algorithms.

Applications

We mentioned several possible applications of the methods developed. In all those fields there exists a large literature, only a rudimentary impression can be given. *Finding characteristic terms* of a document collection is studied, for example, in Information Retrieval. “Global” term weights are, e.g., Poisson overestimation, dis-

crimination value, and inverse document frequency [MRS08]. Specialized methods have been developed for the related problem of automated keyword extraction [RECC10]. *Terminology extraction* is studied, for example, in [PL05]. Methods for *automated query expansion* are discussed, e.g., in [LSM12]. In corpus linguistics and lexicology, the problem of *finding syntagmatic and paradigmatic relations* is a central topic. [SP93, Sto10] present contributions to finding lexical-semantic relations, collocation, word maps, and paradigmatics of a lexical unit. In all cases, the methods suggested above are unique by using information on the subphrase structure of sentences inferred from general structural properties of the corpus. For achieving optimal results it would be interesting to combine distinct approaches.

7 Conclusion

In this paper we described results of an ongoing research project where we analyze corpora in a completely language independent and unsupervised way without any prior linguistic knowledge. We introduced an algorithm for detecting function words and phrases. In many cases the phrases (combinations of phrases) obtained are linguistic phrases or extensions with “adjuncts” (words connecting linguistic phrases). We then showed how phrases can be split into subphrases and how to determine content words. For Chinese, detection of phrases yields satisfactory results, as to word segmentation it is difficult to detect the best level of phrase decomposition.

A weakness of the main algorithm for partitioning sentences into phrases presented in Sections 3.1-3.5 is that only strings can qualify as phrases that occur within distinct contexts in the corpus. All sequences with just one occurrence in the corpus cannot become phrases. To avoid this problem, and to come closer to phrases in a proper linguistic sense we need to *abstract* from plain symbol sequences in some way. The extensions of the main algorithm presented in Section 3.6 represent one promising path for abstraction that deserves further investigation. The detection of phrase patterns such as those in Section 3.6 and Section 5 combined with the pure character-based structure underlying the approach may be also useful for recognising the morphological structure of the language and its dependence on the local phrase context. Some preliminary experiments in this direction showed that our index-structure provides structural indicators that can be used to extract reliable sets of morphemes and to segment words. We assume that unsupervised statistical methods for morphology induction, e.g. [Gol01, Zem08, CL05], can be combined with methods for mining the phrase structure of the corpus, providing another form of abstraction. This might provide a solid alternative to the deep neural network approach taken by [SO15] where embeddings account for the

local context of the words and their morphological properties.

The applications discussed show that the unsupervised analysis of (sub)phrase structure and the two-directional symmetric index used for corpus representation offer many interesting options for finding new solutions to prominent problems in distinct fields.

In this paper, when analyzing corpora we were radical in the sense that we do not use any kind of explicit linguistic knowledge/resource and no knowledge on the nature/functionality of a symbol (blank, hyphen,...). As a matter of fact, from a practical perspective it is interesting to give up this rigid principle. We intend to adapt the methods obtained to “non-raw” texts. For example, text annotated with POS info just represents another kind of interesting source data [Gol10] for “language mining”. The active use of prior knowledge is one point of future work.

Acknowledgements. We express our gratitude to: Max Hadersbeck for the sentence segmentation of the Accountant Corpus; Thomas Müller for providing us with sentence segmentation of the German and English dumps of Wikipedia from May 2014. Wenpeng Yin for the translations from Chinese to English. Stoyan Mihov, Petar Mitankin and Tinko Tinchev for comments and discussions on the ongoing research. Hinrich Schütze for a discussion on function words. The second co-author would like to explicitly acknowledge the technical lead of the first co-author: in the wake of many joint discussions, it was the first co-author who found and implemented the algorithms suggested in this paper. *This research was carried out within the ISALTeC Project, Grant Agreement 625160 of Marie Curie Intra Fellowship within the 7th European Community Framework Programme.*

References

- [AC75] Alfred Aho and Margaret Corasick. Efficient string matching: An aid to bibliographic search. *Communication of the Association for Computing Machinery*, 18(6):333–340, 1975.
- [ACL93] Mikhail Atallah, Danny Chen, and D. Lee. An optimal algorithm for shortest paths on weighted interval and circular-arc graphs with applications. Technical report, Purdue University, 1993.
- [AF13] Giuseppe Attardi and Antonia Fuschetto. Wikipedia Extractor, 2013. http://medialab.di.unipi.it/wiki/Wikipedia_Extractor.

- [BBH⁺87] A. Blumer, J. Blumer, D. Haussler, R. McConnell, and A. Ehrenfeucht. Compete inverted files for efficient text retrieval. *Journal of the Association for Computing Machinery*, 34(3):578–595, July 1987.
- [BTS87] D. E. Ben-Tal, A. abd Brown and R. L. Smith. Posterior Convergence under Incomplete Information. Technical report, Department of Industrial and Operations Engineering, University of Michigan Ann Arbor, 1987.
- [CL05] Mathias Creutz and Krista Lagus. *Unsupervised morpheme segmentation and morphology induction from text corpora using Morfessor 1.0*. Helsinki University of Technology, 2005.
- [CQZH15] Xinchu Chen, Xipeng Qiu, Chenxi Zhu, and Xuanjing Huang. Gated recursive neural network for chinese word segmentation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1744–1753, 2015.
- [CWB⁺11] Ronan Collobert, Jason Weston, Léon Bottou, Micahel Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537, 2011.
- [DIK11] Shweta C. Dharmadhikari, Maya Ingle, and Parag Kulakrni. Empirical studies on machine learning based text classification algorithms. *Advanced Computing: An International Journal*, 2(6):161–169, 2011.
- [DLT04] François Denis, Aurélien Lemay, and Alain Terlutte. Learning regular languages using rfsas. *Theoretical computer science*, 313(2):267–294, 2004.
- [EWN11] Sarah Ebling, Andy Way, Martin Volk, and Sudip Kumar Naskar. Combining semantic and syntactic generalization in example-based machine translation. In *Proceedings of the 15th Conference of the European Association for Machine Translation*, pages 209–216, 2011.
- [Fri52] Charles Fries. *The Structure of English. An Introduction to the Construction of English Sentences*. New York: Harcourt, Brace and Company, 1952.
- [GGJ09] Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. A bayesian framework for word segmentation: Exploring effects of context. *Cognition*, pages 21–54, 2009.

- [GJ08] Marian Grendar and George G. Judge. Consistency of Empirical Likelihood and Maximum A-Posteriori Probability under Misspecification. Technical report, Department of Agricultural & Resource Economics, University of California, Berkley, 2008.
- [Gol01] John A. Goldsmith. Unsupervised learning of the morphology of a natural language. *Computational linguistics*, 2(27):153—189, 2001.
- [Gol10] Johannes Goller. *Exploring the Grammar of Natural Language Using Index Structures*. PhD thesis, Center for language and information processing (CIS), University of Munich, 2010.
- [Gro93] Maurice Gross. Local grammars and their representation by finite automata. In M. Hoey, editor, *Data, Description, Discourse, Papers on the English Language in honour of John McH Sinclair*, pages 26–38. Harper-Collins, London, 1993.
- [Gro97] Maurice Gross. The construction of local grammars. In Emmanuel Roche and Yves Schabes, editors, *Finite-State Language Processing*, pages 329–354. MIT Press, 1997.
- [Gro99] Maurice Gross. A bootstrap method for constructing local grammars. In *Contemporary Mathematics: Proceedings of the Symposium, University of Belgrad*, pages 229–250, Belgrad, 1999.
- [Gue05] Franz Guenther. Corpus Calculus in a nutshell. Technical report, Center for language and information processing (CIS), University of Munich, 205.
- [Gus97] Dan Gusfield. *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge university press, 1997.
- [Har54] Zellig S. Harris. Distributional structure. *Word*, 10:146–162, 1954.
- [Har91] Zellig S. Harris. *A theory of language and information. A mathematical approach*. Oxford University Press, 1991.
- [HNR68] Peter Hart, Nils Nillson, and Bertram Raphael. A formal basis for the heuristic determination of minimal cost paths. *Systems Science and Cybernetics, IEEE Transactions on*, 4(2):100–107, 1968.
- [HNR72] Peter Hart, Nils Nillson, and Bertram Raphael. Correction to ”a formal basis for the heuristic determination of minimum cost paths”.

Newsletter Association for Computational Machinery SIGAI, (37):28–29, 1972.

- [htt] <http://www.ebuk.cn>. Bodyguard legend. <http://dzs.i7wu.cn/files/7/15352/201305041326472943.rar>.
- [HW03] Barbara Höhle and Jürgen Weissenborn. German-learning infants’ ability to detect unstressed closed-class elements in continuous speech. *Developmental Science*, pages 122–127, 2003.
- [IHS⁺01] Shunsuke Inenaga, Hiromasa Hoshino, Ayumi Shinohara, Masayuki Takeda, and Setsuo Arikawa. On-line construction of symmetric compact directed acyclic word graphs. In *Proc. of 8th International Symposium on String Processing and Information Retrieval (SPIRE’01)*, pages 96–110. IEEE Computer Society, 2001.
- [Koe10] Philipp Koehn. *Statistical Machine Translation*. Cambridge University Press, 2010.
- [KS06] Tibor Kiss and Jan Strunk. Unsupervised multilingual sentence boundary detection. *Computational Linguistics*, 2006.
- [LK09] Percy Liang and Dan Klein. Online em for unsupervised models. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, NAACL ’09, pages 611–619, Stroudsburg, PA, USA, 2009. Association for Computational Linguistics.
- [LSM12] Thomas Lüke, Philipp Schaer, and Philipp Mayr. Improving retrieval results with discipline-specific query expansion. In *Theory and Practice of Digital Libraries*, pages 408–413. Springer, 2012.
- [LST13] Hybrid Speech Recognition With Deep Bidirectional LSTM. Graves, alex and jaitly, navdeep and mohamed, abdel-rahman. In *Automatic Speech Recognition and Understanding Workshop, 2013*, pages 273–278, 2013.
- [Lui07] Yuan Lui. Extraction of significant phrases from text. *International Journal of Computer, Control, Quantum and Information Engineering*, 1(9), 2007.
- [MRS08] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.

- [ON04] Franz Och and Hermann Ney. The alignment template approach to statistical machine translation. *Computational Linguistics*, 30(4):417–449, 2004.
- [PFM04] Fuchun Peng, Fangfang Feng, and Andre McCallum. Chinese segmentation and new word detection using conditional random fields. Technical report, University of Massachusetts - Amherst, 2004.
- [PL05] Alexandre Patry and Philippe Langlais. Corpus-based terminology extraction. In *Proceedings of the 7th International Conference on Terminology and Knowledge Engineering*, pages 313–321, 2005.
- [RECC10] Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. Automatic keyword extraction from individual documents. *Text Mining*, pages 1–20, 2010.
- [Sak10] Yasubumi Sakakibara. Probabilistic context-free grammars. *Encyclopedia of Machine Learning*, pages 802–805, 2010.
- [Seb01] Fabrizio Sebastiani. Machine learning in automated text categorization. *ArXiv e-prints*, 2001.
- [SO15] Radu Soricut and Franz Och. Unsupervised morphology induction using word embeddings. In *Proc. NAACL*, 2015.
- [SP93] Hinrich Schütze and Jan Pedersen. A vector model for syntagmatic and paradigmatic relatedness. In *Proc. of Ninth Annual Conference of the UW Centre for the New OED and Text Research*, pages 104–113. Oxford, England, 1993.
- [Sto10] Petra Storjohann. *Lexical-semantic Relations: Theoretical and Practical Perspectives*. *Linguisticae investigationes: Supplementa : LIS : studies in French and general linguistics*. John Benjamins Publishing Company, 2010.
- [SWC06] Rushen Shi, Janet Werker, and Anne Cutler. Recognition and representation of function words in english learning infants. *Infancy*, 10(2):187–198, 2006.
- [Tom13] Masaru Tomita. *Efficient parsing for natural language: A fast algorithm for practical systems*, volume 8. Springer Science & Business Media, 2013.

- [Tra07] Rosemarie Tracy. *Wie Kinder Sprachen Lernen: Und wie wir sie dabei unterstützen können*. Nar Francke Attempto Verlag, 2007.
- [TSM15] Kai Sheng Tai, Richard Socher, and Christopher D. Manning. Improved semantic representations form tree-structured long short-term memory networks. *ArXiv e-prints*, 2015.
- [Ukk95] Esko Ukkonen. On-line construction of suffix-trees. *Algorithmica*, 14(3):249–260, 1995.
- [Zem08] Daniel Zeman. Unsupervised acquiring of morphological paradigms from tokenized text. In *Advances in Multilingual and Multimodal Information*, pages 892–899. Springer Berlin Heidelberg, 2008.
- [ZLC⁺13] Zede Zhu, Miao LI, Lei Chen, Zhenxin Yang, and Sheng Chen. Combination of unsupervised keyphrase extraction algorithms. In *2013 International Conference on Asian Language Processing*, pages 33–36, 2013.

Appendix A

We present top segments of the ranked lists of terminological expressions obtained for some kernels using the method described in Section 5.2. No filtering or cleansing steps were applied to improve the lists.

Example 7.1 Terminological expressions obtained for kernel “disease”:

disease/8418 disease./832 diseases/526 disease./390 diseases./348 diseases./158 this disease./153 coronary artery disease/85 infectious diseases/85 diseased/82 Alzheimer’s disease/82 this disease/77 coronary artery disease./73 liver disease./68 cardiovascular disease./63 liver disease/53 Alzheimer’s disease (AD)/52 Alzheimer’s disease./52 heart disease/52 Parkinson’s disease./51 disease-specific/51 heart disease./50 disease)/47 Crohn’s disease/46 disease-free survival/46 coronary heart disease./45 autoimmune diseases./45 metastatic disease./45 cardiovascular disease/43 Alzheimer’s disease/42 diseases such/42 cardiovascular disease./42 The disease/41 autoimmune diseases/41 vascular disease/40 cardiovascular diseases/40 Alzheimer’s disease (AD)/39 coronary heart disease (CHD)/39 coronary artery disease./38 Parkinson’s disease/37 cardiovascular diseases./37 disease progression/36 coronary heart disease/35 renal disease./34 metastatic disease/33 disease-related/33 liver disease./31 infectious diseases./31 Parkinson’s disease (PD)/31 Crohn’s disease./30 disease-free/30 Parkinson’s disease./30 this disease./29 course of the disease./29 inflammatory bowel disease/29 autoimmune disease/28 coronary heart disease./28 vascular disease./28 diseases, including/28 inflammatory bowel disease./28 coronary artery disease (CAD)/28 Alzheimer’s disease (AD)/28 renal disease/27 disease-/27 Hodgkin’s disease/26 cardiovascular disease (CVD)/26 neurodegenerative diseases/26 Hodgkin’s disease./25 periodontal disease./25 Crohn’s disease./25 neurodegenerative diseases./25 kidney disease./25 inflammatory diseases/25 disease)/24 obstructive pulmonary disease (COPD)/24 diseases, such/23 lung disease./23 Parkinson’s disease (PD)/23 and disease./23 disease severity/23 lung disease/22 periodontal disease/22 liver diseases/22 inflammatory diseases./22 ischemic heart disease/22 disease, including/21 obstructive pulmonary disease (COPD)/21 heart disease./21 coeliac disease./20 ischaemic heart disease./20 disease./20 autoimmune disease./20 coronary artery disease (CAD)/20 Graves’ disease./20 cardiovascular diseases./20 Parkinson’s disease (PD)/20 diseases in/19 congenital heart disease/19

various diseases/19 liver diseases./19 sickle cell disease./19 renal disease./19 sexually transmitted diseases/19 advanced disease./19 end-stage renal disease (ESRD)/19 respiratory diseases/19 diseases associated/19 and diseased/18 vascular diseases/18 inflammatory bowel disease./18 graft-versus-host disease (GVHD)/18 genetic diseases/18 kidney disease/18 kidney disease (CKD)/18 chronic obstructive pulmonary disease (COPD)/18 lung diseases/17 underlying disease./17 Paget's disease/17 allergic diseases/17 stable disease./17 cerebrovascular disease./17 infectious disease/17 disease-associated/17 end-stage renal disease./17 cerebrovascular disease./17 pulmonary disease./17 infectious disease./17 malignant diseases/16 ischemic heart disease./16 disease, especially/16 malignant disease./16 disease, particularly/16 coronary disease./16 progressive disease/16 other diseases/16 kidney disease./16 disease activity/16 course of the disease/15 vascular disease./15 a disease/15 Hodgkin's disease./15 cardiac disease/15 Graves' disease/15 congenital heart disease./15 periodontal diseases/15 chronic diseases/15 infectious diseases./15 recurrent disease./15 the disease/15 severity of the disease./15 inflammatory bowel disease (IBD)/15 coronary heart disease (CHD)/15 various diseases./15 obstructive pulmonary disease./15 disease, and/15 recurrent disease/15

Example 7.2 Terminological expressions obtained for kernel “syndrome”:

syndrome/549 syndrome./287 syndrome./173 syndromes/157 metabolic syndrome/100 syndromes./82 this syndrome/62 syndromes./60 metabolic syndrome./47 nephrotic syndrome/44 this syndrome./33 acquired immunodeficiency syndrome (AIDS)/33 metabolic syndrome./31 Down syndrome/30 syndrome)/27 Down's syndrome/26 Sjögren's syndrome/26 The syndrome/22 irritable bowel syndrome/20 syndrome)/20 syndrome and/20 carpal tunnel syndrome/19 Cushing's syndrome/19 syndrome characterized/19 respiratory distress syndrome/17 a syndrome/16 acquired immune deficiency syndrome (AIDS)/16 Cushing's syndrome./16 compartment syndrome/16 This syndrome/15 irritable bowel syndrome (IBS)/14 long QT syndrome/13 Down syndrome./13 toxic shock syndrome/13 nephrotic syndrome./12 Cushing's syndrome./12 Rett syndrome/12 myelodysplastic syndrome (MDS)/12 meconium aspiration syndrome/11 clinical syndrome/11 syndrome associated/11 -syndrome/11 Budd-Chiari syndrome/11 Turner's syndrome/11 Down syndrome./11 coronary syndromes/11 polycystic ovary syndrome (PCOS)/11 syndrome"/11 Tourette syndrome/11 this syndrome./10 Down's syndrome./10 syndrome is/10 The metabolic syndrome/10 antiphospholipid syndrome/10 metabolic syndrome (MetS)/10 's syndrome/9 syndrome is presented./9 Wolff-Parkinson-White syndrome/9 Down syndrome (DS)/9 respiratory distress syndrome./9 acute coronary syndrome/9 hepatopulmonary syndrome/9 acquired immunodeficiency syndrome/9 pain syndromes/8 Marfan's syndrome/8 adult respiratory distress syndrome/8 respiratory distress syndrome (RDS)/8 nephrotic syndrome./8 Reiter's syndrome/8 syndrome in/8 syndrome, including/8 Turner syndrome/8 Tourette's syndrome./8 Marfan syndrome/8 myelodysplastic syndrome/8 coronary syndrome/8 Metabolic syndrome/8 pain syndrome/8 pain syndrome./8 Marfan syndrome./8 Lennox-Gastaut syndrome/8 sicca syndrome/8 Cockayne syndrome/8 Sheehan's syndrome/7 Goodpasture's syndrome/7 Sjögren's syndrome./7 respiratory distress syndrome./7 pain syndromes./7 acquired immunodeficiency syndrome (AIDS)./7 syndrome is characterized/7 Felty's syndrome/7 polycystic ovary syndrome/7 withdrawal syndrome/7 sleep apnoea syndrome/7 Terson syndrome/7 coronary syndromes./7 Williams syndrome/7 acquired immunodeficiency syndrome./7 coronary syndrome./7 fatigue syndrome (CFS)/7 Rett syndrome./7 Behçet's syndrome/7 SVC syndrome/6 Bartter's syndrome/6) syndrome/6 Zollinger-Ellison syndrome/6 Down's syndrome./6 syndrome occurred/6 adult respiratory distress syndrome (ARDS)/6 Tourette's syndrome/6 Wiskott-Aldrich syndrome/6 fragile X syndrome./6 obstructive sleep apnea syndrome/6 sudden infant death syndrome (SIDS)/6 Guillain-Barré syndrome/6 compartment syndrome./6 syndrome"./6 obstructive sleep apnea syndrome (OSAS)/6 syndromes such/6 respiratory distress syndrome (ARDS)/6 systemic inflammatory response syndrome/6 wasting syndrome/6 Lowe syndrome/6 metabolic syndrome (MS)/6 Brugada syndrome/6 Apert syndrome/6 Beckwith-Wiedemann syndrome/6 Ogilvie's syndrome/6 syndrome patients./6 fragile X-associated tremor/ataxia syndrome/6 Gilbert's syndrome/5 syndrome due/5 's syndrome./5 deficiency syndrome/5 fat embolism syndrome/5 -like syndrome/5 syndrome (P/5 pain syndromes./5 Reye's syndrome/5 adult respiratory distress syndrome./5 syndrome'/5 sick sinus syndrome/5 abstinence syndrome/5 Raynaud's syndrome/5 acquired immune deficiency syndrome (AIDS)/5 myelodysplastic syndromes/5 heart-asthenia syndrome/5 Sjögren's syndrome./5 Marfan's syndrome./5 syndrome caused/5 syndromes, including/5 HELLP syndrome/5 The clinical syndrome/5 carpal tunnel syndrome./5 sudden infant death syndrome (SIDS)/5 short bowel syndrome./5 syndrome (/5 post-polio syndrome/5 Turner syndrome./5 antiphospholipid syndrome./5 Zollinger-Ellison syndrome (ZES)/5 fragile X syndrome./5 sudden infant death syndrome/5 Guillain-Barré syndrome./5 syndrome (HUS)/5 syndrome are de-

scribed./5 Angelman syndrome/5 Kallmann syndrome/5 systemic inflammatory response syndrome (SIRS)/5 acquired immune deficiency syndrome/5 syndrome occurs/5 Stevens-Johnson syndrome/5 Klinefelter's syndrome/5 Bartter syndrome/5 the metabolic syndrome/5 Panayiotopoulos syndrome/5 Noonan syndrome/5 polycystic ovary syndrome (PCOS)/5 polycystic ovary syndrome./5 Sweet's syndrome/5 metabolic syndrome and/5 acute coronary syndromes (ACS)/5 syndrome was diagnosed/5 syndrome was diagnosed./5 syndrome-associated/5 WAGR syndrome/5 Churg-Strauss syndrome (CSS)/5 Sjogren's syndrome/5 post-Q-fever fatigue syndrome/5 coronary syndrome./5 coronary syndromes (ACS/5 Hunter syndrome./5 Eisenmenger syndrome./5 syndrome, especially/5 scimitar syndrome/5 5q- syndrome/5

Example 7.3 Terminological expressions obtained for kernel “inflammatory”:

inflammatory/381 anti-inflammatory/117 inflammatory response/111 proinflammatory/63 proinflammatory cytokines/52 inflammatory response./46 antiinflammatory/41 pro-inflammatory/40 inflammatory reaction/37 inflammatory process/35 inflammatory cells/34 inflammatory mediators/33 inflammatory cytokines/33 inflammatory bowel/32 anti-inflammatory effects/31 inflammatory responses/31 pro-inflammatory cytokines/30 inflammatory bowel disease/29 inflammatory bowel disease./28 inflammatory cells./28 inflammatory markers/27 an inflammatory/26 inflammatory diseases/25 inflammatory lesions/24 proinflammatory cytokine/24 pro-inflammatory cytokine/23 inflammatory processes/23 chronic inflammatory/23 anti-inflammatory effect/22 inflammatory diseases./22 inflammatory processes./21 inflammatory responses./21 inflammatory changes/21 inflammatory response./19 inflammatory process./19 inflammatory bowel disease./18 inflammatory./17 inflammatory reaction/17 inflammatory reactions/17 anti-inflammatory drugs/16 inflammatory disorders/15 inflammatory conditions./15 inflammatory bowel disease (IBD)/15 the inflammatory/15 inflammatory infiltrate/14 inflammatory cell infiltration/14 non-steroidal anti-inflammatory drugs (NSAIDs)/14 systemic inflammatory response/14 nonsteroidal anti-inflammatory drugs/13 inflammatory disease./13 nonsteroidal anti-inflammatory drugs (NSAIDs)/13 inflammatory conditions/13 anti-inflammatory activity/12 inflammatory cytokine/12 anti-inflammatory properties/12 proinflammatory mediators/12 anti-inflammatory activity./12 inflammatory markers./12 The inflammatory response/11 inflammatory mediators./11 systemic inflammatory/11 inflammatory bowel disease (IBD)/11 inflammatory markers./11 inflammatory bowel disease (IBD)/11 The inflammatory/10 inflammatory disease/10 proinflammatory cytokines./10 inflammatory process./10 anti-inflammatory agents/9 inflammatory diseases./9 inflammatory disorders./9 inflammatory bowel diseases/9 The anti-inflammatory activity/9 inflammatory cytokines./9 inflammatory cytokines./9 pro- and anti-inflammatory/9 pelvic inflammatory/8 anti-inflammatory drugs./8 inflammatory demyelinating/8 noninflammatory/8 anti-inflammatory properties./8 anti-inflammatory effects./8 proinflammatory cytokines./8 neuroinflammatory/8 inflammatory cell/7 inflammatory changes./7 inflammatory infiltrate./7 inflammatory mediator/7 The anti-inflammatory/7 inflammatory parameters/7 inflammatory diseases such/7 infiltration of inflammatory/7 nonsteroidal anti-inflammatory drug/7 inflammatory reactions./7 inflammatory cascade/7 anti-inflammatory cytokine IL-10/7 inflammatory disease./7 inflammatory state/7 inflammatory activity/6 nonsteroidal anti-inflammatory drugs./6 inflammatory infiltrates/6 non-inflammatory/6 Anti-inflammatory/6 various inflammatory/6 anti-inflammatory cytokine/6 inflammatory pseudotumor/6 inflammatory pain/6 inflammatory events/6 inflammatory conditions./6 inflammatory infiltration/6 inflammatory mediators such/6 systemic inflammatory response syndrome/6 inflammatory mediators./6 pelvic inflammatory disease./6 non-steroidal anti-inflammatory drug (NSAID)/6 macrophage inflammatory protein-1alpha/6 of inflammatory/6 inflammatory states./6 pro-inflammatory cytokines./6 pro-inflammatory cytokines./6 Pro-inflammatory cytokines/6 inflammatory responses./6 non-steroidal anti-inflammatory drugs/6 macrophage inflammatory/6 inflammatory changes./6 Proinflammatory cytokines/6 anti-inflammatory agents./5 fibro-inflammatory/5 pelvic inflammatory disease (PID)/5 inflammatory pain./5 nonsteroidal antiinflammatory/5 inflammatory reactions./5 necroinflammatory/5 inflammatory and neoplastic/5 inflammatory skin/5 inflammatory lung/5 An inflammatory/5 proinflammatory mediators./5 non-specific inflammatory/5 systemic inflammatory response syndrome (SIRS)/5 inflammatory reaction./5 proinflammatory cytokine./5 active inflammatory bowel/5 nonsteroidal anti-inflammatory drugs./5 inflammatory disorder/5 inflammatory stimuli/5 postinflammatory/5 non-steroidal anti-inflammatory drugs./5 production of proinflammatory/5 inflammatory component./5 systemic inflammatory response./5 several inflammatory/5 inflammatory status/5 inflammatory cell infiltration./5 inflammatory signaling/5 anti-inflammatory activities./5 inflammatory stimuli./5 non-steroidal anti-inflammatory drug/5 nonsteroidal antiinflammatory drugs (NSAIDs)/5 release of inflammatory mediators/5